a. Domain Name Service (DNS) encompassing

Microsoft DNS

From Wikipedia, the free encyclopedia
Jump to: navigation, search

Microsoft DNS is the name given to the implementation of domain name system services provided in Microsoft Windows operating systems.

Contents

[hide]

[edit] Overview

The Domain Name System support in Microsoft Windows NT, and thus its derivatives Windows 2000, Windows XP, and Windows Server 2003, comprises two clients and a server. Every Microsoft Windows machine has a DNS lookup client, to perform ordinary DNS lookups. Some machines have a Dynamic DNS client, to perform Dynamic DNS Update transactions, registering the machines' names and IP addresses. Some machines run a DNS server, to publish DNS data, to service DNS lookup requests from DNS lookup clients, and to service DNS update requests from DNS update clients.

The server software is only supplied with the server versions of Windows.

[edit] DNS lookup client

Applications perform DNS lookups with the aid of a DLL. They call library functions in the DLL, which in turn handle all communications with DNS servers (over UDP or TCP) and return the final results of the lookup back to the applications.

Microsoft's DNS client also has optional support for local caching, in the form of a *DNS Client* service (also known as *DNSCACHE*). Before they attempt to directly communicate with DNS servers, the library routines first attempt to make a local IPC connection to the DNS Client service on the machine. If there is one, and if such a connection can be made, they hand the actual work of dealing with the lookup over to the DNS Client service. The DNS Client service itself communicates with DNS servers, and caches the results that it receives.

Microsoft's DNS client is capable of talking to multiple DNS servers. The exact algorithm varies according to the version, and service pack level, of the operating system; but in general all communication is with a *preferred* DNS server until it fails to answer, whereupon communication switches to one of several *alternative* DNS servers.

[edit] The effects of running the DNS Client service

There are several minor differences in system behavior depending on whether the DNS Client service is started:

- **Parsing of the "hosts" file**: The lookup functions read only the hosts file if they cannot off-load their task onto the DNS Client service and have to fall back to communicating with DNS servers themselves. In turn, the DNS Client service reads the "hosts" file once, at startup, and only re-reads it if it notices that the last modification timestamp of the file has changed since it last read it. Thus:
    - With the DNS Client service running: The "hosts" file is read and parsed only a few times, once at service startup, and thereafter whenever the DNS Client service notices that it has been modified.
    - Without the DNS Client service running: The "hosts" file is read and parsed repeatedly, by each individual application program as it makes a DNS lookup.

- **The effect of multiple answers in the "hosts" file**: The DNS Client service does not use the "hosts" file directly when performing lookups. Instead, it (initially) populates its cache from it, and then performs lookups using the data in its cache. When the lookup functions fall back to doing the work themselves, however, they scan the "hosts" file directly and sequentially, stopping when the first answer is found. Thus:
    - With the DNS Client service running: If the "hosts" file contains multiple lines denoting multiple answers for a given lookup, all of the answers in the cache will be returned.
    - Without the DNS Client service running: If the "hosts" file contains multiple lines denoting multiple answers for a given lookup, only the first answer found will be returned.

- **Fallback from preferred to alternative DNS servers**: The fallback from the preferred DNS server to the alternative DNS servers is done by whatever entity, the DNS Client service or the library functions themselves, is actually performing the communication with them. Thus:
    - With the DNS Client service running: Fallback to the alternative DNS servers happens globally. If the preferred DNS server fails to answer, all subsequent communication is with the alternative DNS servers.

  - Without the DNS Client service running: Any fallback to the alternative DNS servers happen locally, within each individual process that is making DNS queries. Different processes may be in different states, some talking to the preferred DNS server and some talking to alternative DNS servers.

[edit] Differences from other systems

Linux distributions and various versions of Unix have a generalized name resolver layer. The resolver can be controlled to use a **hosts** file or Network Information Service (NIS), by configuring the Name Service Switch.

[edit] Dynamic DNS Update client

Whilst DNS lookups read DNS data, DNS updates *write* them. Both workstations and servers running Windows attempt to send Dynamic DNS update requests to DNS servers.

Workstations running Windows attempt to register their names and their IP addresses with DNS servers, so that other machines may locate them by name. Prior to Windows Vista (and Windows Server 2008) this registration is performed by the *DHCP Client* service. It is thus necessary to run the DHCP Client service on pre-Vista machines, even if DHCP isn't being used to configure the machine in order to dynamically register a machine's name and address for DNS lookup. The DHCP Client service registers name and address data whenever they are changed (either manually by an administrator or automatically by the granting or revocation of a DHCP lease). In Microsoft Vista (and Windows Server 2008) Microsoft moved the registration functionallity from the *DHCP Client* service to the *DNS Client* service.

Servers running Microsoft Windows also attempt to register other information, in addition to their names and IP addresses, such as the locations of the LDAP and Kerberos services that they provide.

[edit] DNS server

Microsoft Windows server operating systems can run the *DNS Server* service. This is a monolithic DNS server that provides many types of DNS service, including caching, Dynamic DNS update, zone transfer, and DNS notification. DNS notification implements a push mechanism for notifying a select set of secondary servers for a zone when it is updated.

Microsoft's "DNS Server" service was first introduced in Windows NT 3.51 as an add-on with Microsoft's collection of BackOffice services (at the time was marked to be used for testing proposes only). Its code is a fork of ISC's BIND, version 4.3, and remains largely compatible with it including the format of all master files.

As of 2004, it was the fourth most popular DNS server (counting BIND version 9 separately from versions 8 and 4) for the publication of DNS data.[1]

Like various other DNS servers, Microsoft's DNS server supports different database *back ends*. Microsoft's DNS server supports two such back ends. DNS data can be stored either in *master files* (also known as *zone files*) or in the Active Directory database itself. In the latter case, since Active Directory (rather than the DNS server) handles the actual replication of the database across multiple machines, the database can be modified on any server ("multiple-master replication"), and the addition or removal of a *zone* will be immediately propagated to all other DNS servers within the appropriate Active Directory "replication scope". (Contrast this with BIND, where when such changes are made, the list of *zones*, in the `/etc/named.conf` file, has to be explicitly updated on each individual server.)

Microsoft's DNS server can be administered using either a graphical user interface, the "DNS Management Console", or a command line interface, the **dnscmd** utility.

[edit] Common issues

Prior to Windows Server 2003 and Microsoft Windows 2000 Service Pack 3, the most common problem encountered with Microsoft's DNS server was cache pollution. Although Microsoft's DNS Server had a mechanism for properly dealing with cache pollution, the mechanism was turned off by default. [2]

A longstanding well-known issue is incompatibility with BIND configuration files, in particular, the lack of support for DNS wildcards. This can be partially attributed to the fact that Microsoft's DNS server is based on BIND 4.3, before BIND added the support for DNS wildcards. *Loose Wildcarding* can be enabled. To create a wildcard character record, the *Dnscmd* command-line tool can be used. Also the support of IPv6 is implemented using a different technique from that of BIND 9, further driving even more incompatibilities between the two products.

In 2004, a common problem involved the feature of the Windows Server 2003 version of Microsoft's DNS server to use EDNS0, which a large number of firewalls could not cope with. [3]

[edit] See also

Dynamic DNS

From Wikipedia, the free encyclopedia
Jump to: navigation, search

Dynamic DNS is a method / protocol / network service that provides the capability for a networked device, such as a router or computer system using the Internet Protocol Suite, to notify a Domain Name System (DNS) name server to change, in real time, the active DNS configuration of its configured hostnames, addresses or other information.

A popular application of dynamic DNS is to provide a residential user's Internet gateway that has a variable, often changing, IP address with a well known hostname resolvable through standard DNS queries.

Contents

[hide]

[edit] History

In the initial stages of the Internet (ARPANET) addressing of hosts on the network was achieved by static translation tables that mapped hostnames to IP addresses. The tables were maintained manually in form of the host file. The Domain Name System brought a method of distributing the same address information automatically online through recursive queries to remote databases configured for each network, or domain. Even this DNS facility still used static lookup tables at each participating node. IP addresses, once assigned to a particular host, rarely changed and the mechanism was initially sufficient. However, the rapid growth of the Internet and the proliferation of personal computers in the workplace and in homes created the substantial burden for administrators of keeping track of assigned IP addresses and managing their address space. The Dynamic Host Configuration Protocol (DHCP) allowed enterprises and Internet service providers (ISPs) to assign addresses to computers automatically as they powered up. In addition, this helped conserve the address space available, since not all devices might be actively used at all times and addresses could be assigned as needed. This feature required that DNS servers be kept current automatically as well. The first implementations of *dynamic DNS* fulfilled this purpose: Host computers gained the feature to notify their respective DNS server of the address they had received from a DHCP server or through self-configuration. This protocol-based DNS update method was documented and standardized in IETF publication RFC 2136 in 1997 and has become a standard part of the DNS protocol (see also nsupdate program).

The explosive growth and proliferation of the Internet into people's homes brought a growing shortage of available IP addresses. DHCP became an important tool for ISPs as well to manage their address spaces for connecting home and small-business end-users with a single IP address each by implementing network address translation (NAT) at the customer premise router. The private network behind these routers uses address space set aside for these purposes (RFC 1918), masqueraded by the NAT device. This, however, broke the end-to-end principle of Internet architecture and methods were required to allow private networks, with frequently changing external IP addresses, to discover their public address and insert it into the Domain Name System in order to participate in Internet communications more fully. Today, numerous providers, called *Dynamic DNS* service providers, offer such technology and services on the Internet.

[edit] Function

Dynamic DNS providers offer a software client program that automates the discovery and registration of client's public IP addresses. The client program is executed on a computer or device in the private network. It connects to the service provider's systems and causes those systems to link the discovered public IP address of the home network with a hostname in the domain name system. Depending on the provider, the hostname is registered within a domain owned by the provider or the customer's own domain name. These services can function by a number of mechanisms. Often they use an HTTP service request since even restrictive environments usually allow HTTP service. This group of services is commonly also referred to by the term *Dynamic DNS*, although it is not the standards-based DNS Update method. However, the latter might be involved in the providers systems.

Most home networking routers today have this feature already built into their firmware. One of the early routers to support Dynamic DNS was the UMAX UGate-3000 in 1999, which supported the TZO.COM dynamic DNS service.[1]

An example is residential users who wish to access their personal computer at home while traveling. If the home computer has a fixed static IP address, the user can connect directly using this address, but many provider networks force frequent changes the IP address configured in their customers' equipment. With dynamic DNS, the home computer can automatically associate its current IP address with a domain name. As a result the remote user can resolve the host name used for the dynamic DNS service entry to the current address of the home computer with a DNS query. If a remote control program such as VNC server may be kept running on a host in the private network, the user can connect to the home network with a VNC client program.

In Microsoft Windows networks, dynamic DNS is an integral part of Active Directory, because domain controllers register their network service types in DNS so that other computers in the Domain (or Forest) can access them.

Increasing efforts to secure Internet communications today involve encryption of all dynamic updates via the public Internet, as these public dynamic DNS services have been abused increasingly to design security breaches. Standards-based methods within the DNSSEC protocol suite, such as TSIG, have been developed to secure DNS updates, but are not widely in use. Microsoft developed alternative technology (GSS-TSIG) based on Kerberos authentication.

[edit] See also

- Domain name system
- DNS hosting service
- Name server
- Fast-flux DNS
- Comparison of DNS server software

[edit] External links

- Dynamic DNS Services at the Open Directory Project
- Understanding Dynamic Update in Windows Server 2008 Help

- Domain Name System on Microsoft TechNet
- How to configure

Root nameserver

From Wikipedia, the free encyclopedia
Jump to: navigation, search

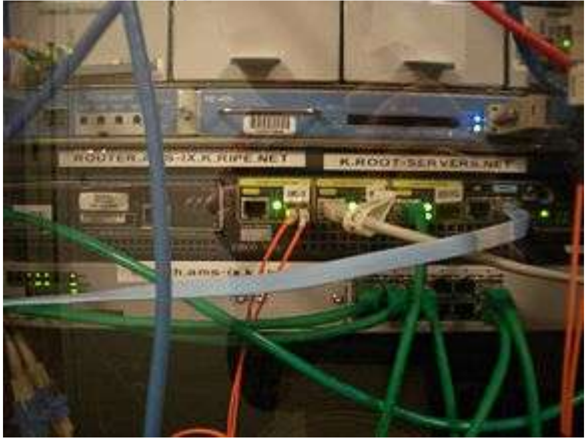

A Cisco 7301 router, part of the AMS-IX mirror of the K root-server.

A **root name server** is a name server for the Domain Name System's root zone. It directly answers requests for records in the root zone and answers other requests returning a list of the designated authoritative name servers for the appropriate top-level domain (TLD). The root name servers are a critical part of the Internet because they are the first step in translating (resolving) human readable host names into IP addresses that are used in communication between Internet hosts.

A combination of limits in the DNS and certain protocols, namely the practical size of unfragmented User Datagram Protocol (UDP) packets, resulted in a limited number of root server addresses that can be accommodated in DNS name query responses. This limit has determined the number of name server installations at (currently) 13 clusters, serving the needs of the entire public Internet worldwide.

Contents

[hide]

[edit] Root domain

The Domain Name System is a hierarchical naming system for computers, services, or any resource participating in the Internet. The top of that hierarchy is the root domain. The root domain does not have a formal name and its label in the DNS hierarchy is an empty string. All fully qualified domain names (FQDNs) on the Internet can be regarded as ending with this empty string for the root domain, and therefore ending in a full stop character (the label delimiter), e.g., `www.example.com.`. This is generally implied rather than explicit, as modern DNS software does not actually require that the terminating dot be included when attempting to translate a domain name to an IP address.

The root domain contains all top-level domains of the Internet. As of June 2009, there are 20 generic top-level domains (gTLDs) and 248 country code top-level domains (ccTLDs) in the root domain.[1] In addition, the `ARPA` domain is used for technical name spaces in the management of Internet addressing and other resources. A `TEST` domain is used for testing internationalized domain names.

[edit] Resolver operation

When a computer on the Internet needs to resolve a domain name, it uses resolver software to perform the lookup. A resolver breaks the name up into its labels from right to left. The first component (TLD) is requeried using a root server to obtain the responsible authoritative server. Queries for each label return more specific name servers until a name server returns the answer of the original query.

In practice, most of this information does not change very often over a period of hours and therefore it is cached by intermediate name servers or by a name cache built into the user's application. DNS lookups to the root nameservers may therefore be relatively infrequent. A survey in 2003 [2] reports that only 2% of all queries to the root servers were legitimate. Incorrect or non-existent caching was responsible for 75% of the queries, 12.5% were for unknown TLDs, 7% were for lookups using IP addresses as if they were domain names, etc. Some misconfigured desktop computers even tried to update the root server records for the TLDs. A similar list of observed problems and recommended fixes has been published in RFC 4697.

Although any local implementation of DNS can implement its own private root name servers, the term "root name server" is generally used to describe the thirteen well-known root name servers that implement the root name space domain for the Internet's official global implementation of the Domain Name System.

[edit] Root server addresses

There are currently 13 root name servers specified, with names in the form *`letter`*`.root-servers.net`, where *`letter`* ranges from A to M. This does not mean there are 13 physical servers; each operator uses redundant computer equipment to provide reliable service even if failure of hardware or software occur. Additionally, nine of the servers operate in multiple geographical locations using a routing technique called anycast, providing increased performance and even more fault tolerance.

Ten servers were originally in the United States; some are now operated via anycast. Three servers were originally located in Stockholm (I), Amsterdam (K), and Tokyo (M).

Letter	**IPv4 address**	**IPv6 address**	**Old name**	**Operator**	**Location**	**Software**
A	198.41.0.4	2001:503:ba3e::2:30	ns.internic.net	VeriSign	distributed using anycast	BIND
B	192.228.79.201 (since January 2004; originally was 128.9.0.107)[3]	2001:478:65::53 (not in root zone yet)	ns1.isi.edu	USC-ISI	Marina Del Rey, California, U.S.	BIND
C	192.33.4.12		c.psi.net	Cogent Communications	distributed using anycast	BIND
D	128.8.10.90		terp.umd.edu	University of Maryland	College Park, Maryland, U.S.	BIND
E	192.203.230.10		ns.nasa.gov	NASA	Mountain View, California, U.S.	BIND
F	192.5.5.241	2001:500:2f::f	ns.isc.org	Internet Systems Consortium	distributed using anycast	BIND 9[4]
G	192.112.36.4		ns.nic.ddn.mil	Defense Information Systems Agency	distributed using anycast	BIND
H	128.63.2.53	2001:500:1::803f:235	aos.arl.army.mil	U.S. Army Research Lab	Aberdeen Proving Ground, Maryland, U.S.	NSD
I	192.36.148.17	2001:7fe::53	nic.nordu.net	Autonomica	distributed using anycast	BIND
J	192.58.128.30 (since	2001:503:c27::2:30		VeriSign	distributed using	BIND

	November 2002; originally was 198.41.0.10)				anycast	
K	193.0.14.129	2001:7fd::1		RIPE NCC	distributed using anycast	NSD[5]
L	199.7.83.42 (since November 2007; originally was 198.32.64.12)[6]	2001:500:3::42		ICANN	distributed using anycast	NSD[7]
M	202.12.27.33	2001:dc3::35		WIDE Project	distributed using anycast	BIND

Older servers had their own name before the policy of using similar names was established.

The choice of 13 nameservers was made because of limitations in the original DNS specification, which specifies a maximum packet size of 512 bytes when using the User Datagram Protocol (UDP).[8] The addition of IPv6 addresses for the root nameservers requires more than 512 bytes, which is facilitated by the EDNS0 extension to the DNS standard.[9] While only 13 names are used for the root nameservers, there are many more physical servers; C, F, I, J, K, L and M servers now exist in multiple locations on different continents, using anycast address announcements to provide decentralized service. As a result most of the physical root servers are now outside the United States, allowing for high performance worldwide.

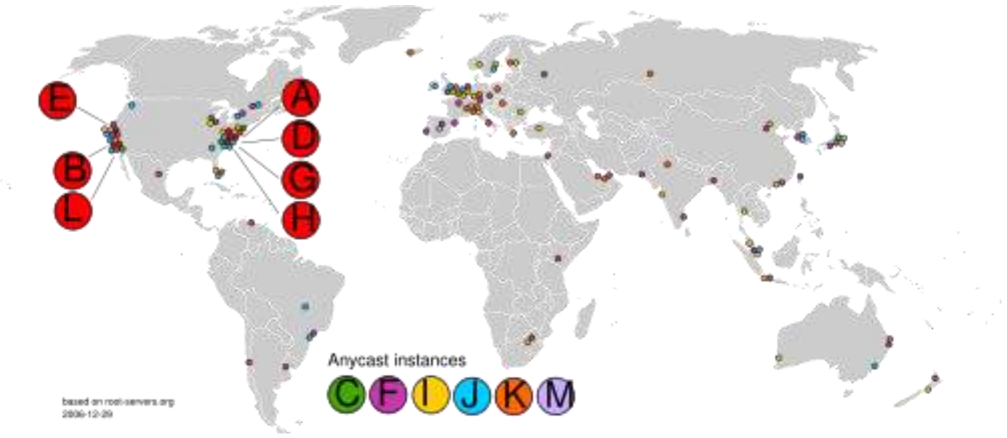

At the end of 2006 there were a total of 13 root nameservers, including Anycast servers.

There are also several alternative namespace systems with an alternative DNS root using their own set of root nameservers that exist in parallel to the mainstream nameservers. The first, AlterNIC, generated a substantial amount of press.[citation needed]

The function of a root name server may also be implemented locally, or on a provider network. Such servers are synchronized with the official root zone file as published by ICANN, and do not constitute an alternate root.

As the root nameservers are an important part of the Internet, they have come under attack several times, although none of the attacks have ever been serious enough to severely affect the performance of the Internet.

[edit] Root server supervision

The DNS Root Server System Advisory Committee is an ICANN committee. ICANN's bylaws[10] assign authority over the operation of the root nameservers of the Domain Name System to the DNS Root Server System Advisory Committee.

[edit] Root zone file

The root zone file is a small (about 200 kB) data set whose publication is the primary purpose of root nameservers.

The root zone file is at the apex of a hierarchical distributed database called the Domain Name System (DNS). This database is used by almost all Internet applications to translate worldwide unique names like www.wikipedia.org into other identifiers such as IP addresses.

The contents of the root zone file is a list of names and numeric IP addresses of the authoritative DNS servers for all top-level domains (TLDs) such as com, org, edu, or the country code top-level domains. On 12 December 2004, there were 258 TLDs and 773 different authoritative servers for those TLDs listed. Other name servers forward queries for which they do not have any information about authoritative servers to a root name server. The root name server, using its root zone file, answers with a referral to the authoritative servers for the appropriate TLD or with an indication that no such TLD exists.[11]

[edit] See also

- Distributed denial of service attacks on root nameservers
- EDNS0 (Extended DNS, version 0)
- Internet backbone
- Open Root Server Network
- Blackhole server

[edit] References

1. ^ "List of top-level domains". ICANN. http://data.iana.org/TLD/tlds-alpha-by-domain.txt.
2. ^ "Wow, That's a Lot of Packets" (PDF). 2003. http://dns.measurement-factory.com/writings/wessels-pam2003-paper.pdf. Retrieved 2008-02-05.
3. ^ "New IPv4 address for b.root-servers.net". http://www.root-servers.org/news/new-ip-b.html.
4. ^ F-root | Internet Systems Consortium

5. ^ K-root Homepage
6. ^ "Advisory — "L Root" changing IP address on 1st November". ICANN. http://blog.icann.org/?p=227.
7. ^ l.root-servers.net
8. ^ RFC 1035 Domain names - implementation and specification
9. ^ ICANN: Accommodating IP Version 6 Address Resource Records for the Root of the Domain Name System
10. ^ ICANN Bylaws XI-2.3
11. ^ ISOC, *DNS Root Name Servers explained for the non-expert*, (Available online, accessed 19 March 2010.)

Notes

- Root Server Technical Operations Association
- Root Servers' Geographical

Installing and Configuring Zones

The following introduction provides high-level planning information for global and non-global zones. For overview and planning information and specific procedures, see Chapter 16, Introduction to Solaris Zones, in System Administration Guide: Solaris Containers-Resource Management and Solaris Zones.

Solaris Zones Partitioning Technology (Overview)

After the Solaris OS is installed, you can install and configure zones. The global zone is the single instance of the operating system that is running and is contained on every Solaris system. The global zone is both the default zone for the system and the zone that is used for system-wide administrative control. A non-global zone is a virtualized operating system environment.

Solaris Zones are a software partitioning technology used to virtualize operating system services and provide an isolated and secure environment for running applications. When you create a zone, you produce an application execution environment in which processes are isolated from all other zones. This isolation prevents processes that are running in one zone from monitoring or affecting processes that are running in any other zones. Even a process running in a non-global zone with superuser credentials cannot view or affect activity in any other zones. A process running in the global zone with superuser credentials can affect any process in any zone.

Understanding Global and Non-Global Zones

The global zone is the only zone from which a non-global zone can be configured, installed, managed, or uninstalled. Only the global zone is bootable from the system hardware. Administration of the system infrastructure, such as physical devices, routing, or dynamic reconfiguration (DR), is only possible in the global zone. Appropriately privileged processes running in the global zone can access objects associated with any or all other zones. The following table summarizes the characteristics of both global and non-global zones.

Global Zone	**Non-Global Zone**
Is assigned ID 0 by the system	Is assigned a zone ID by the system when the zone is booted
Provides the single instance of the Solaris kernel that is bootable and running on the system	Shares operation under the Solaris kernel booted from the global zone
Contains a complete installation of the Solaris system software packages	Contains an installed subset of the complete Solaris Operating System software packages
Can contain additional software packages or additional software, directories, files, and other data not installed through packages	Contains Solaris software packages shared from the global zone
Provides a complete and consistent product database that contains information about all software components installed in the global zone	Can contain additional installed software packages not shared from the global zone Can contain additional software, directories, files, and other data created on the non-global zone that are not installed through packages or shared from the global zone
Holds configuration information specific to the global zone only, such as the global zone host name and file system table	Has configuration information specific to that non-global zone only, such as the non-global zone host name and file system table
Is the only zone that is aware of all devices and all file systems	Has a complete and consistent product database that contains information about all software components installed on the zone, whether present on the non-global zone or shared read-only from the global zone
Is the only zone with knowledge of non-global zone existence and configuration	Is not aware of the existence of any other zones
Is the only zone from which a non-global zone can be configured, installed, managed, or uninstalled	Cannot install, manage, or uninstall other zones, including itself

For more information, see the following:

- Chapter 16, Introduction to Solaris Zones, in System Administration Guide: Solaris Containers-Resource Management and Solaris Zones
- Installing and Configuring Zones

Solaris Zones (Planning)

After the Solaris OS is installed, you can install and configure zones. The global zone is the single instance of the operating system that is running and is contained on every Solaris system. The global zone is both the default zone for the system and the zone that is used for system-wide administrative control. A non-global zone is a virtualized operating system environment.

Caution –

Any command that accepts an alternate root (/) file system by using the **-R** option or equivalent must not be used if the following are true:

- The command is run in the global zone.
- The alternative root (/) file system refers to any path within a non-global zone.

An example is the **-R** *root_path* option to the pkgadd utility run from the global zone with a path to the root (/) file system in a non-global zone.

For a list of utilities that accept an alternate root (/) file system and more information about zones, see Restriction on Accessing A Non-Global Zone From the Global Zone in System Administration Guide: Solaris Containers-Resource Management and Solaris Zones.

Installing and Upgrading When Using Non-global Zones

When the Solaris OS is installed, the software group installed in the global zone is the set of packages that is shared by all the non-global zones. For example, if you install the Entire software group, all zones contain these packages. By default, any additional packages installed in the global zone also populate the non-global zones. You can segregate into non-global zones any applications, namespaces, servers, and network connections such as NFS and DHCP as well as other software. Each non-global zone is unaware of other non-global zones and each can operate independently. For example, you might have installed the Entire software group on the global zone and have running on separate non-global zones the Java Enterprise System Messaging Server, a database, DHCP, and a web server. When installing non-global zones remember the performance requirements of the applications running in each non-global zone.

Caution –

A Solaris Flash archive cannot be properly created when a non-global zone is installed. The Solaris Flash feature is not compatible with Solaris Zones partitioning technology. If you create a Solaris Flash archive, the resulting archive is not installed properly when the archive is deployed under these conditions:

- The archive is created In a non-global zone
- The archive is created in a global zone that has non-global zones installed

---

Upgrading When Non-Global Zones Are Installed

Starting with the Solaris 10 1/06 release, when you are upgrading the Solaris OS, you can upgrade a system that has non-global zones installed. The Solaris interactive installation program and custom JumpStart programs enable an upgrade.

- With the Solaris interactive installation program, you can upgrade a system with non-global zones by selecting the Upgrade Install on the Select Upgrade or Initial Install panel. The installation program then analyzes your system to determine if your system is upgradable, and provides you a summary of the analysis. The installation program then prompts you to continue the upgrade. You can use this program with the following limitations:

  You cannot customize your upgrade. For example, you cannot install additional software products, install additional locale packages, or modify the disk layout.

  You must use the Solaris Operating System DVD or a DVD-created network installation image. You cannot use the Solaris Software CDs to upgrade a system. For more information about installing with this program, see Chapter 2, Installing With the Solaris Installation Program (Tasks), in Solaris 10 Installation Guide: Basic Installations.
- With the custom JumpStart installation program, you can upgrade by using only the `install_type` and `root_device` keywords.

  Because some keywords affect non-global zones, some keywords cannot be included in a profile. For example, using keywords that add packages, reallocate disk space, or add locales would affect non-global zones. If you use these keywords, they are ignored or cause the JumpStart upgrade to fail. For a list of these keywords, see Limiting Profile Keywords When Upgrading With Non-Global Zones.

Caution –

You cannot use Solaris Live Upgrade to upgrade a system when non-global zones are installed. You can create a boot environment with the `lucreate` command, but if you use the `luupgrade` command, the upgrade fails. An error message is displayed.

Disk Space Requirements for Non-Global Zones

When installing the global zone, be sure to reserve enough disk space for all of the zones you might create. Each non-global zone might have unique disk space requirements. The following description is a brief overview of planning information. For complete planning requirements and recommendations, see Chapter 18, Planning and Configuring Non-Global Zones (Tasks), in System Administration Guide: Solaris Containers-Resource Management and Solaris Zones.

No limits are placed on how much disk space can be consumed by a zone. The global zone administrator is responsible for space restriction. Even a small uniprocessor system can support a number of zones running simultaneously.

The characteristics of the packages installed in the global zone affect the space requirements of the non-global zones that are created. The number of packages and space requirements are factors. The following are general disk space guidelines.

- Approximately 100 Mbytes of free disk space is suggested when the global zone has been installed with all of the standard Solaris packages. Increase this amount if additional packages are installed in the global zone. By default, any additional packages installed in the global zone also populate the non-global zones. The directory location in the non-global zone for these additional packages is specified through the `inherit-pkg-dir` resource.
- Add 40 Mbytes of RAM per zone if the system has sufficient swap space. This addition is recommended to make each zone operational. When planning your system size, consider this addition of RAM.

Platform Names and Groups

When you are adding clients for a network installation, you must know your system architecture (platform group). If you are writing a custom JumpStart installation rules file, you need to know the platform name.

Some examples of platform names and groups follow. For a full list of SPARC based systems, see Solaris Sun Hardware Platform Guide at `http://docs.sun.com/`.

Table 3–8 Example of Platform Names and Groups

System	Platform Name	Platform Group
Sun Fire	T2000	sun4v
Sun Blade^{TM}	SUNW,Sun-Blade-100	sun4u

System	Platform Name	Platform Group
x86 based	i86pc	i86pc

Note –

On a running system, you can also use the `uname -i` command to determine a system's **platform name** or the `uname -m` command to determine a system's **platform group**.

SPARC: 64–bit Packaging Changes

In previous Solaris releases, the Solaris OS was delivered in separate packages for 32-bit and 64-bit components. In the Solaris 10 OS, packaging has been simplified with the delivery of most 32-bit and 64-bit components in a single package. The combined packages retain the names of the original 32-bit packages, and the 64-bit packages are no longer delivered. This change reduces the number of packages and simplifies installation. This change means that you might need to modify your custom JumpStart script or other package installation scripts to remove references to the 64-bit packages.

The 64-bit packages are renamed with the following conventions:

- If a 64-bit package has a 32-bit counterpart, the 64-bit package is named with the 32-bit package name. For example, a 64-bit library such as /usr/lib/sparcv9/libc.so.1 previously would have been delivered in SUNWcslx, but now is delivered in SUNWcsl. The 64-bit SUNWcslx package is no longer delivered.
- If a package does not have a 32-bit counterpart, the “x” suffix is removed from the name. For example, SUNW1394x becomes SUNW1394.

- *Previous*: Installing and Configuring Zones
- *Next*: x86: Partitioning Recommendations

x86: Partitioning Recommendations

When using the Solaris OS on x86 based systems, follow these guidelines for partitioning your system.

The Solaris installation program uses a default boot-disk partition layout. These partitions are called `fdisk` partitions. An fdisk partition is a logical partition of a disk drive that is dedicated to a particular operating system on x86 based systems. To install the Solaris software, you must set up at least one Solaris `fdisk` partition on an x86 based system. x86 based systems allow up to four different `fdisk` partitions on a disk. These partitions can be used to hold individual operating systems. Each operating system must be located on a unique `fdisk` partition. A system can only have one Solaris `fdisk` partition per disk.

Table 3–9 x86: Default Partitions

Partitions	**Partition Name**	**Partition Size**
First partition (on some systems)	Diagnostic or Service partition	Existing size on system.
Second partition (on some systems)	x86 boot partition	**For the Solaris 10 3/05 release**: An x86 boot partition is created and is the existing size on the system. **Starting with the Solaris 10 1/06 release**, the following conditions apply: • If you are performing an initial installation, this partition is not created. • If you upgrade and your system does not have an existing x86 boot partition, this partition is not created. • If you upgrade and your system has an x86 boot partition: If the partition is required to bootstrap from one boot device to another, the x86 boot partition is preserved on the system. If the partition is not required to boot additional boot devices, the x86 boot partition is removed. The contents of the partition are moved to the root partition.
Third partition	Solaris OS partition	Remaining space on the boot disk.

Default Boot-Disk Partition Layout Preserves the Service Partition

The Solaris installation program uses a default boot-disk partition layout to accommodate the diagnostic or Service partition. If your system currently includes a diagnostic or Service partition, the default boot-disk partition layout enables you to preserve this partition.

Note –

If you install the Solaris OS on an x86 based system that does not currently include a diagnostic or Service partition, the installation program does not create a new diagnostic or Service partition by default. If you want to create a diagnostic or Service partition on your system, see your hardware documentation.

Chapter 4 x86: GRUB Based Booting For Solaris Installation

This chapter describes the GRUB based booting on x86 based systems that relates to Solaris installation. This chapter contains the following sections:

- x86: GRUB Based Booting (Overview)
- x86: GRUB Based Booting (Planning)
- x86: Locating the GRUB Menu's `menu.lst` File (Tasks)

x86: GRUB Based Booting (Overview)

Starting with the Solaris 10 1/06 release, GRUB, the open source boot loader, has been adopted as the default boot loader in the Solaris OS.

Note –

GRUB based booting is not available on SPARC based systems.

The **boot loader** is the first software program that runs after you power on a system. After you power on an x86 based system, the Basic Input/Output System (BIOS) initializes the CPU, the memory, and the platform hardware. When the initialization phase has completed, the BIOS loads the boot loader from the configured boot device, and then transfers control of the system to the boot loader.

GRUB is an open source boot loader with a simple menu interface that includes boot options that are predefined in a configuration file. GRUB also has a command-line interface that is accessible from the menu interface for performing various boot commands. In the Solaris OS, the GRUB implementation is compliant with the Multiboot Specification. The specification is described in detail at `http://www.gnu.org/software/grub/grub.html`.

Because the Solaris kernel is fully compliant with the Multiboot Specification, you can boot a Solaris x86 based system by using GRUB. With GRUB, you can more easily boot and install various operating systems. For example, on one system, you could individually boot the following operating systems:

- Solaris OS

- Microsoft Windows

Note –

GRUB detects Microsoft Window partitions but does not verify that the OS can be booted.

A key benefit of GRUB is that it is intuitive about file systems and kernel executable formats, which enables you to load an operating system without recording the physical position of the kernel on the disk. With GRUB based booting, the kernel is loaded by specifying its file name, and the drive, and the partition where the kernel resides. GRUB based booting replaces the Solaris Device Configuration Assistant and simplifies the booting process with a GRUB menu.

x86: How GRUB Based Booting Works

After GRUB gains control of the system, a menu is displayed on the console. In the GRUB menu, you can do the following:

- Select an entry to boot your system
- Modify a boot entry by using the built-in GRUB edit menu
- Manually load an OS kernel from the command line

A configurable timeout is available to boot the default OS entry. Pressing any key aborts the default OS entry boot.

To view an example of a GRUB menu, see Description of the GRUB Main Menu.

x86: GRUB Device Naming Conventions

The device naming conventions that GRUB uses are slightly different from previous Solaris OS versions. Understanding the GRUB device naming conventions can assist you in correctly specifying drive and partition information when you configure GRUB on your system.

The following table describes the GRUB device naming conventions.

Table 4–1 Naming Conventions for GRUB Devices

Device Name	Description
(`fd0`), (`fd1`)	First diskette, second diskette

Device Name	Description
(`nd`)	Network device
(`hd0,0`), (`hd0,1`)	First and second `fdisk` partition of first `bios` disk
`(hd0,0,a`), (`hd0,0,b`)	Solaris/BSD slice 0 and 1 on first `fdisk` partition on the first `bios` disk

Note –

All GRUB device names must be enclosed in parentheses. Partition numbers are counted from 0 (zero), not from 1.

For more information about `fdisk` partitions, see Guidelines for Creating an fdisk Partition in System Administration Guide: Devices and File Systems.

x86: Where to Find Information About GRUB Based Installations

For more information about these changes, see the following references:

Table 4–2 Where to Find Information on GRUB Based Installations

Topic	GRUB Menu Tasks	For More Information
Installation	To install from the Solaris OS CD or DVD media	Solaris 10 Installation Guide: Basic Installations.
	To install from a network installation image	Part II, Installing Over a Local Area Network, in Solaris 10 Installation Guide: Network-Based Installations
	To configure a DHCP server for network installations	Preconfiguring System Configuration Information With the DHCP Service (Tasks) in Solaris 10 Installation Guide: Network-Based Installations
	To install with the Custom JumpStart	x86: Performing a Custom JumpStart Installation

Topic	GRUB Menu Tasks	For More Information
	program	
	To activate or fall back to a boot environment by using Solaris Live Upgrade	• Activating a Boot Environment in Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning • Chapter 10, Failure Recovery: Falling Back to the Original Boot Environment (Tasks), in Solaris 10 Installation Guide: Solaris Live Upgrade and Upgrade Planning
System Administration	For more detailed information about GRUB and for administrative tasks	Chapter 11, GRUB Based Booting (Tasks), in System Administration Guide: Basic Administration

x86: GRUB Based Booting (Planning)

This section describes the basics of GRUB based booting and describes the GRUB menu.

When you install the Solaris OS, two GRUB menu entries are installed on the system by default. The first entry is the Solaris OS entry. The second entry is the failsafe boot archive, which is to be used for system recovery. The Solaris GRUB menu entries are installed and updated automatically as part of the Solaris software installation and upgrade process. These entries are directly managed by the OS and should not be manually edited.

During a standard Solaris OS installation, GRUB is installed on the Solaris `fdisk` partition without modifying the system BIOS setting. If the OS is not on the BIOS boot disk, you need to do one of the following:

- Modify the BIOS setting.
- Use a boot manager to bootstrap to the Solaris partition. For more details, see your boot manager.

The preferred method is to install the Solaris OS on the boot disk. If multiple operating systems are installed on the machine, you can add entries to the `menu.lst` file. These entries are then displayed in the GRUB menu the next time you boot the system.

For additional information on multiple operating systems, see How Multiple Operating Systems Are Supported in the GRUB Boot Environment in System Administration Guide: Basic Administration.

x86: Performing a GRUB Based Installation From the Network

Performing a GRUB based network boot requires a DHCP server that is configured for PXE clients and an install server that provides `tftp` service. The DHCP server must be able to respond to the DHCP classes, `PXEClient` and `GRUBClient`. The DHCP response must contain the following information:

- IP address of the file server
- Name of the boot file (`pxegrub`)

Note –

`rpc.bootparamd`, which is usually a requirement on the server side for performing a network boot, is not required for a GRUB based network boot.

If no PXE or DHCP server is available, you can load GRUB from CD-ROM or local disk. You can then manually configure the network in GRUB and download the multiboot program and the boot archive from the file server.

For more information, see Overview of Booting and Installing Over the Network With PXE in Solaris 10 Installation Guide: Network-Based Installations.

Description of the GRUB Main Menu

When you boot an x86 based system, the GRUB menu is displayed. This menu provides a list of boot entries to choose from. A **boot entry** is an OS instance that is installed on your system. The GRUB menu is based on the `menu.lst` file, which is a configuration file. The `menu.lst` file is created by the Solaris installation program and can be modified after installation. The `menu.lst` file dictates the list of OS instances that are shown in the GRUB menu.

- If you install or upgrade the Solaris OS, the GRUB menu is automatically updated. The Solaris OS is then displayed as a new boot entry.
- If you install an OS other than the Solaris OS, you must modify the `menu.lst` configuration file to include the new OS instance. Adding the new OS instance enables the new boot entry to appear in the GRUB menu the next time that you boot the system.

Example 4–1 GRUB Main Menu

In the following example, the GRUB main menu shows the Solaris and Microsoft Windows operating systems. A Solaris Live Upgrade boot environment is also listed that is named `second_disk`. See the following for descriptions of each menu item.

```
GNU GRUB version 0.95 (616K lower / 4127168K upper memory)
+-------------------------------------------------------------------+
|Solaris                                                            |
|Solaris failsafe                                                   |
|second_disk                                                        |
|second_disk failsafe                                               |
|Windows                                                            |
+-------------------------------------------------------------------+
Use the ^ and v keys to select which entry is highlighted. Press
enter to boot the selected OS, 'e' to edit the commands before
booting, or 'c' for a command-line.
```

Solaris

Specifies the Solaris OS.

Solaris failsafe

Specifies a boot archive that can be used for recovery if the Solaris OS is damaged.

second_disk

Specifies a Solaris Live Upgrade boot environment. The `second_disk` boot environment was created as a copy of the Solaris OS. It was upgraded and activated with the `luactivate` command. The boot environment is available for booting.

Windows

Specifies the Microsoft Windows OS. GRUB detects these partitions but does not verify that the OS can be booted.

Description of GRUB `menu.lst` File

The GRUB `menu.lst` file lists the contents of the GRUB main menu. The GRUB main menu lists boot entries for all the OS instances that are installed on your system, including Solaris Live Upgrade boot environments. The Solaris software upgrade process preserves any changes that you make to this file.

Any revisions made to the `menu.lst` file are displayed on the GRUB main menu, along with the Solaris Live Upgrade entries. Any changes that you make to the file become effective at the next system reboot. You can revise this file for the following reasons:

- To add to the GRUB menu entries for operating systems other than Solaris

- To customize booting behavior such as specifying the default OS on the GRUB menu

---

⚠ Caution –

Do not use the GRUB `menu.lst` file to modify Solaris Live Upgrade entries. Modifications could cause Solaris Live Upgrade to fail.

Although you can use the `menu.lst` file to customize booting behavior such as booting with the kernel debugger, the preferred method for customization is to use the `eeprom` command. If you use the `menu.lst` file to customize, the Solaris OS entries might be modified during a software upgrade. Changes to the file would then be lost.

For information about how to use the `eeprom` command, see How to Set Solaris Boot Parameters by Using the eeprom Command in System Administration Guide: Basic Administration.

Example 4–2 `Menu.lst` File

Here is a sample of a `menu.lst` file:

```
default 0
timeout 10
title Solaris
  root (hd0,0,a)
  kernel /platform/i86pc/multiboot -B console=ttya
  module /platform/i86pc/boot_archive
title Solaris failsafe
  root (hd0,0,a)
  kernel /boot/multiboot -B console=ttya -s
  module /boot/x86.miniroot.safe
#----- second_disk - ADDED BY LIVE UPGRADE - DO NOT EDIT  -----
title second_disk
  root (hd0,1,a)
  kernel /platform/i86pc/multiboot
  module /platform/i86pc/boot_archive
title second_disk failsafe
  root (hd0,1,a)
  kernel /boot/multiboot kernel/unix -s
  module /boot/x86.miniroot-safe
#----- second_disk -------------- END LIVE UPGRADE ------------
title Windows
  root (hd0,0)
  chainloader -1
```

`default`

Specifies which item to boot if the timeout expires. To change the default, you can specify another item in the list by changing the number. The count begins with zero for the first title. For example, change the default to 2 to boot automatically to the `second_disk` boot environment.

`timeout`

Specifies the number of seconds to wait for user input before booting the default entry. If no timeout is specified, you are required to choose an entry.

`title` ***OS name***

Specifies the name of the operating system.

- If this is a Solaris Live Upgrade boot environment, *OS name* is the name you gave the new boot environment when it was created. In the previous example, the Solaris Live Upgrade boot environment is named `second_disk`.
- If this is a failsafe boot archive, this boot archive is used for recovery when the primary OS is damaged. In the previous example, Solaris failsafe and `second_disk` failsafe are the recovery boot archives for the Solaris and `second_disk` operating systems.

`root (hd0,0,a)`

Specifies on which disk, partition, and slice to load files. GRUB automatically detects the file system type.

`kernel /platform/i86pc/multiboot`

Specifies the multiboot program. The kernel command must always be followed by the multiboot program. The string after multiboot is passed to the Solaris OS without interpretation.

For a complete description of multiple operating systems, see How Multiple Operating Systems Are Supported in the GRUB Boot Environment in System Administration Guide: Basic Administration.

Locating the `menu.lst` File To Change the GRUB Menu

You must always use the `bootadm` command to locate the GRUB menu's `menu.lst` file. The `list-menu` subcommand finds the active GRUB menu. The `menu.lst` file lists all the operating systems that are installed on a system. The contents of this file dictate the list of operating systems that is displayed on the GRUB menu. If you want to make changes to this file, see x86: Locating the GRUB Menu's `menu.lst` File (Tasks).

x86: Locating the GRUB Menu's `menu.lst` File (Tasks)

Starting with the Solaris 10 1/06 release, the GRUB menu can be updated. For example, you might want to change the default time for how fast the default OS is booted. Or, you might want to add another OS to the GRUB menu.

Typically, the active GRUB menu's `menu.lst` file is located at `/boot/grub/menu.lst`. In some situations, the GRUB `menu.lst` file resides elsewhere. For example, in a system that uses Solaris Live Upgrade, the GRUB `menu.lst` file might be on a boot environment that is not the currently running boot environment. Or if you have upgraded a system with an x86 boot partition, the `menu.lst` file might reside in the `/stubboot` directory. Only the active GRUB `menu.lst` file is used to boot the system. In order to modify the GRUB menu that is displayed when you boot the system, the active GRUB `menu.lst` file must be modified. Changing any other GRUB `menu.lst` file has no effect on the menu that is displayed when you boot the system. To determine the location of the active GRUB `menu.lst` file, use the `bootadm` command. The `list-menu` subcommand displays the location of the active GRUB menu. The following procedures determine the location of the GRUB menu's `menu.lst` file.

For more information about the `bootadm` command, see bootadm(1M) man page.

Locating the GRUB Menu's `menu.lst` file

In the following procedure, the system contains two operating systems: Solaris and a Solaris Live Upgrade boot environment, `second_disk`. The Solaris OS has been booted and contains the GRUB menu.

Steps

1. Become superuser or assume an equivalent role.

   Roles contain authorizations and privileged commands. For more information about roles, see Configuring RBAC (Task Map) in System Administration Guide: Security Services.

2. To locate the `menu.lst` file, type:

   ```
   # /sbin/bootadm list-menu
   ```

3. The location and contents of the file are displayed.

```
The location for the active GRUB menu is: /boot/grub/menu.lst
default 0
timeout 10
0 Solaris
1 Solaris failsafe
2 second_disk
3 second_disk failsafe
```

Locating the GRUB Menu's `menu.lst` File When the active `menu.lst` file is in Another Boot Environment

In the following procedure, the system contains two operating systems: `Solaris` and a Solaris Live Upgrade boot environment, `second_disk`. In this example, the `menu.lst` file does not exist in the currently running boot environment. The `second_disk` boot environment has been booted. The `Solaris` boot environment contains the GRUB menu. The `Solaris` boot environment is not mounted.

Steps

1. Become superuser or assume an equivalent role.

   Roles contain authorizations and privileged commands. For more information about roles, see Configuring RBAC (Task Map) in System Administration Guide: Security Services.

2. To locate the `menu.lst` file, type:

```
# /sbin/bootadm list-menu
```

3. The location and contents of the file are displayed.

```
The location for the active GRUB menu is: /dev/dsk/device_name(not mounted)
The filesystem type of the menu device is <ufs>
default 0
timeout 10
0 Solaris
1 Solaris failsafe
2 second_disk
3 second_disk failsafe
```

4. Because the file system containing the `menu.lst` file is not mounted, mount the file system. Specify the UFS file system and the device name.

```
# /usr/sbin/mount -F ufs /dev/dsk/device_name /mnt
```

5. Where *device_name* specifies the location of the root (`/`) file system on the disk device of the boot environment that you want to mount. The device name is entered in the form of `/dev/dsk/c`*w*`t`*x*`d`*y*`s`*z*. For example:

```
# /usr/sbin/mount -F ufs /dev/dsk/c0t1d0s0 /mnt
```

6. You can access the GRUB menu at `/mnt/boot/grub/menu.lst`
7. Unmount the filesystem

```
# /usr/sbin/umount /mnt
```

8.
9. **Note –**
10. If you mount a boot environment or a file system of a boot environment, ensure that the file system or file systems are unmounted after use. If these file systems are not unmounted, future Solaris Live Upgrade operations on that boot environment might fail.
11.

Locating the GRUB Menu's `menu.lst` File When a Solaris Live Upgrade Boot Environment is Mounted

In the following procedure, the system contains two operating systems: `Solaris` and a Solaris Live Upgrade boot environment, `second_disk`. The `second_disk` boot environment has been booted. The `Solaris` boot environment contains the GRUB menu. The `Solaris` boot environment is mounted at `/.alt.Solaris`.

Steps

1. Become superuser or assume an equivalent role.

   Roles contain authorizations and privileged commands. For more information about roles, see Configuring RBAC (Task Map) in System Administration Guide: Security Services.

2. To locate the `menu.lst` file, type:

```
# /sbin/bootadm list-menu
```

3. The location and contents of the file are displayed.

```
The location for the active GRUB menu is:
/.alt.Solaris/boot/grub/menu.lst
default 0
timeout 10
0 Solaris
1 Solaris failsafe
2 second_disk
3 second_disk failsafe
```

4. Since the boot environment containing the GRUB menu is already mounted, then you can access the `menu.lst` file at `/.alt.Solaris/boot/grub/menu.lst.`

Locating the GRUB Menu's `menu.lst` File When Your System Has an x86 Boot Partition

In the following procedure, the system contains two operating systems: Solaris and a Solaris Live Upgrade boot environment, `second_disk`. The `second_disk` boot environment has been booted. Your system has been upgraded and an x86 boot partition remains. The boot partition is mounted at `/stubboot` and contains the GRUB menu. For an explanation of x86 boot partitions, see x86: Partitioning Recommendations.

Steps

1. Become superuser or assume an equivalent role.

   Roles contain authorizations and privileged commands. For more information about roles, see Configuring RBAC (Task Map) in System Administration Guide: Security Services.

2. To locate the `menu.lst` file, type:

```
# /sbin/bootadm list-menu
```

3. The location and contents of the file are displayed.

```
The location for the active GRUB menu is:
/stubboot/boot/grub/menu.lst
default 0
timeout 10
0 Solaris
1 Solaris failsafe
```

```
2 second_disk
3 second_disk failsafe
```

4. You can access the `menu.lst` file at `/stubboot/boot/grub/menu.lst.`

Part II Using Custom JumpStart

This part provides instructions for creating, preparing, and performing custom JumpStart installations.

Contents

Chapter 5 Custom JumpStart (Overview)

This chapter provides an introduction and overview to the custom JumpStart installation process.

Custom JumpStart Introduction

The custom JumpStart installation method is a command–line interface that enables you to automatically install or upgrade several systems, based on profiles that you create. The profiles define specific software installation requirements. You can also incorporate shell scripts to include preinstallation and postinstallation tasks. You choose which profile and scripts to use for installation or upgrade. The custom JumpStart installation method installs or upgrades the system, based on the profile and scripts that you select. Also, you can use a sysidcfg file to specify configuration information so that the custom JumpStart installation is completely hands-off.

Custom JumpStart Example Scenario

The custom JumpStart process can be described by using an example scenario. In this example scenario, the systems need to be set up with the following parameters:

- Install Solaris on 100 new systems.
- Seventy of the systems are SPARC based systems that are owned by the engineering group and need to be installed as standalone systems with the Solaris OS software group for developers.
- The remaining 30 systems are x86 based, owned by the marketing group and need to be installed as standalone systems with the Solaris OS software group for end users.

First, the system administrator must create a rules file and a profile for each group of systems. The rules file is a text file that contains a rule for each group of systems or single systems on which you want to install the Solaris software. Each rule distinguishes a group of systems that are based on one or more system attributes. Each rule also links each group to a profile.

A profile is a text file that defines how the Solaris software is to be installed on each system in the group. Both the rules file and profile must be located in a JumpStart directory.

For the example scenario, the system administrator creates a rules file that contains two different rules, one for the engineering group and another for the marketing group. For each rule, the system's network number is used to distinguish the engineering group from the marketing group.

Each rule also contains a link to an appropriate profile. For example, in the rule for the engineering group, a link is added to the profile, eng_profile, which was created for the engineering group. In the rule for the marketing group, a link is added to the profile, market_profile, which was created for the marketing group.

You can save the rules file and the profiles on a diskette or on a server.

- A profile diskette is required when you want to perform custom JumpStart installations on nonnetworked, standalone systems.
- A profile server is used when you want to perform custom JumpStart installations on networked systems that have access to a server.

After creating the rules file and profiles, validate the files with the check script. If the check script runs successfully, the rules.ok file is created. The rules.ok is a generated version of the rules file that the JumpStart program uses to install the Solaris software.

How the JumpStart Program Installs Solaris Software

After you validate the `rules` file and the profiles, you can begin a custom JumpStart installation. The JumpStart program reads the `rules.ok` file. Then, the JumpStart program searches for the first rule with defined system attributes that match the system on which the JumpStart program is attempting to install the Solaris software. If a match occurs, the JumpStart program uses the profile that is specified in the rule to install the Solaris software on the system.

Figure 5–1 illustrates how a custom JumpStart installation works on a standalone, nonnetworked system. The system administrator initiates the custom JumpStart installation on Pete's system. The JumpStart program accesses the rules files on the diskette in the system's diskette drive. The JumpStart program matches `rule 2` to the system. `rule 2` specifies that the JumpStart program use `Pete's profile` to install the Solaris software. The JumpStart program reads `Pete's profile` and installs the Solaris software, based on the instructions that the system administrator specified in `Pete's profile`.

Figure 5–1 How a Custom JumpStart Installation Works: nonnetworked Example

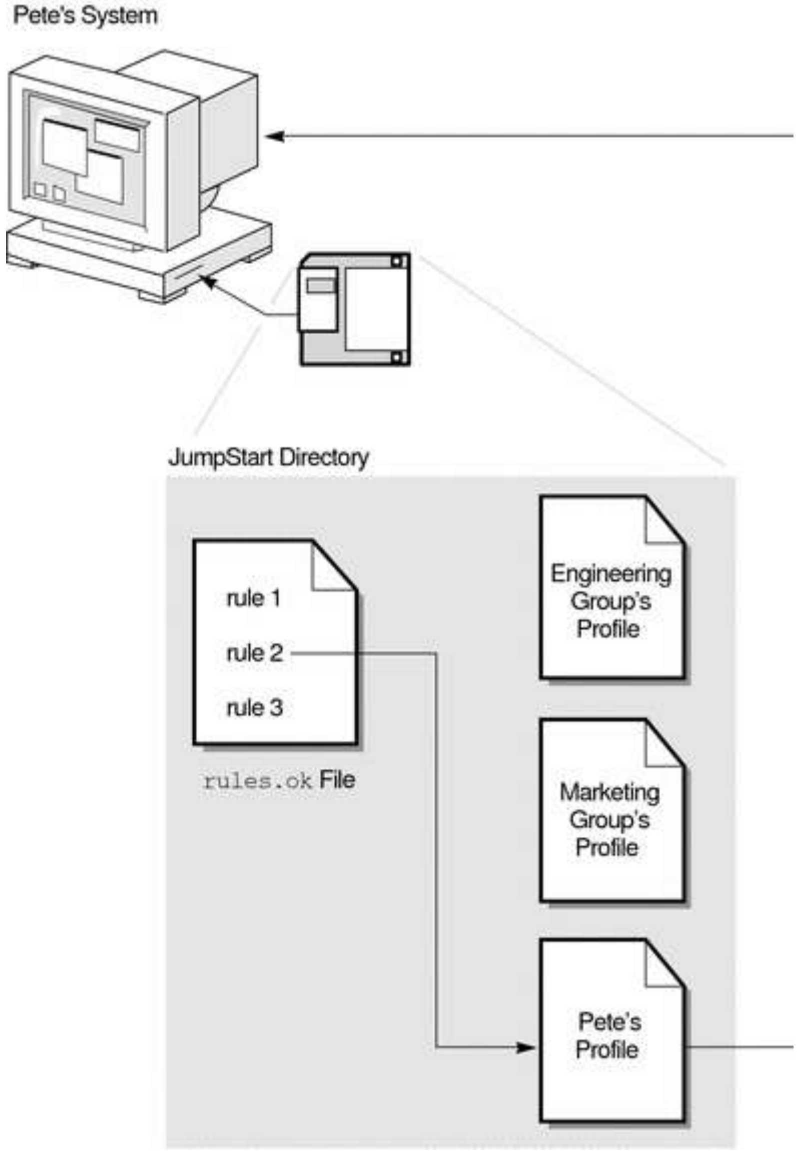

Figure 5–2 illustrates how a custom JumpStart installation works with more than one system on a network. Previously, the system administrator set up different profiles and saved the profiles on a single server. The system administrator initiates the custom JumpStart installation on one of the engineering systems. The JumpStart program accesses the rules files in the `JumpStart/` directory on the server. The JumpStart program matches the engineering system to `rule 1`. `rule 1` specifies that the JumpStart program use `Engineering Group's Profile` to install the Solaris software. The JumpStart program reads `Engineering Group's Profile` and installs the Solaris software, based on the instructions that the system administrator specified in `Engineering Group's Profile`.

Figure 5–2 How a Custom JumpStart Installation Works: Networked Example

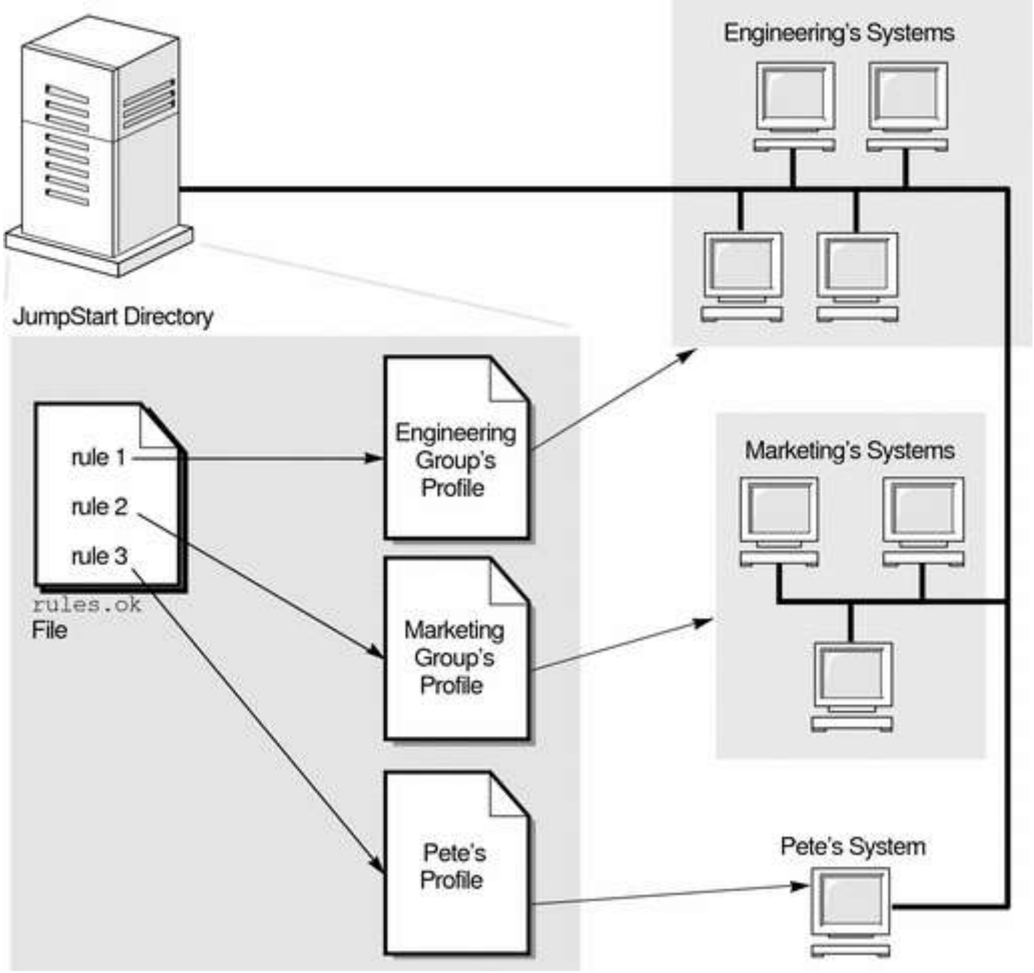

Figure 5–3 describes the order in which the JumpStart program searches for custom JumpStart files.

Figure 5–3 What Happens During a Custom JumpStart Installation

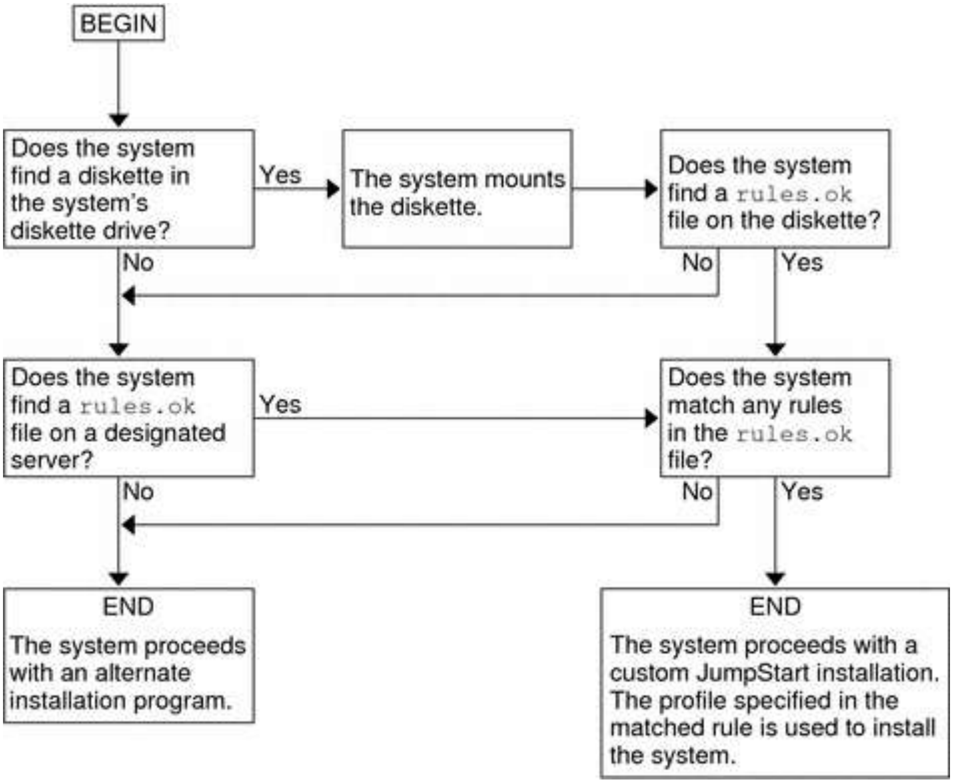
BEGIN
Does the system find a diskette in the system's diskette drive?
Yes
No
The system mounts the diskette.
Does the system find a rules.ok file on the diskette?
No
Yes
Does the system find a rules.ok file on a designated server?
Yes
No
Does the system match any rules in the rules.ok file?
No
Yes
END
The system proceeds with an alternate installation program.
END
The system proceeds with a custom JumpStart installation. The profile specified in the matched rule is used to install the system.

Caching-only servers

A caching-only server saves data in a cache file until the data expires. Expiration occurs based on a ``time-to-live'' field attached to data received from another server.

A caching-only server answers data from its cache if it has the information, or requests it from authoritative servers if it does not.

Caching-only servers may make use of forwarders, servers with connections to outside networks that build up a large database of data. In this case, the caching-only server first queries a forwarder (rather than following its default resolution process), which in turn queries another server, if necessary, to find the desired data. This limits traffic outside the zone to traffic to and from the forwarders. If the forwarder does not answer, the caching-only server tries the primary server for its zone.

Root servers are used by the caching-only server to help find other zones.

Microsoft DNS

From Wikipedia, the free encyclopedia
Jump to: navigation, search

Microsoft DNS is the name given to the implementation of domain name system services provided in Microsoft Windows operating systems.

Contents

[hide]

[edit] Overview

The Domain Name System support in Microsoft Windows NT, and thus its derivatives Windows 2000, Windows XP, and Windows Server 2003, comprises two clients and a server. Every Microsoft Windows machine has a DNS lookup client, to perform ordinary DNS lookups. Some machines have a Dynamic DNS client, to perform Dynamic DNS Update transactions, registering the machines' names and IP addresses. Some machines run a DNS server, to publish DNS data, to service DNS lookup requests from DNS lookup clients, and to service DNS update requests from DNS update clients.

The server software is only supplied with the server versions of Windows.

[edit] DNS lookup client

Applications perform DNS lookups with the aid of a DLL. They call library functions in the DLL, which in turn handle all communications with DNS servers (over UDP or TCP) and return the final results of the lookup back to the applications.

Microsoft's DNS client also has optional support for local caching, in the form of a *DNS Client* service (also known as *DNSCACHE*). Before they attempt to directly communicate with DNS servers, the library routines first attempt to make a local IPC connection to the DNS Client service on the machine. If there is one, and if such a connection can be made, they hand the actual work of dealing with the lookup over to the DNS Client service. The DNS Client service itself communicates with DNS servers, and caches the results that it receives.

Microsoft's DNS client is capable of talking to multiple DNS servers. The exact algorithm varies according to the version, and service pack level, of the operating system; but in general all communication is with a *preferred* DNS server until it fails to answer, whereupon communication switches to one of several *alternative* DNS servers.

[edit] The effects of running the DNS Client service

There are several minor differences in system behavior depending on whether the DNS Client service is started:

- **Parsing of the "hosts" file**: The lookup functions read only the hosts file if they cannot off-load their task onto the DNS Client service and have to fall back to communicating with DNS servers themselves. In turn, the DNS Client service reads the "hosts" file once, at startup, and only re-reads it if it notices that the last modification timestamp of the file has changed since it last read it. Thus:
    - With the DNS Client service running: The "hosts" file is read and parsed only a few times, once at service startup, and thereafter whenever the DNS Client service notices that it has been modified.
    - Without the DNS Client service running: The "hosts" file is read and parsed repeatedly, by each individual application program as it makes a DNS lookup.

- **The effect of multiple answers in the "hosts" file**: The DNS Client service does not use the "hosts" file directly when performing lookups. Instead, it (initially) populates its cache from it, and then performs lookups using the data in its cache. When the lookup functions fall back to doing the work themselves, however, they scan the "hosts" file directly and sequentially, stopping when the first answer is found. Thus:
    - With the DNS Client service running: If the "hosts" file contains multiple lines denoting multiple answers for a given lookup, all of the answers in the cache will be returned.
    - Without the DNS Client service running: If the "hosts" file contains multiple lines denoting multiple answers for a given lookup, only the first answer found will be returned.

- **Fallback from preferred to alternative DNS servers**: The fallback from the preferred DNS server to the alternative DNS servers is done by whatever entity, the DNS Client service or the library functions themselves, is actually performing the communication with them. Thus:
    - With the DNS Client service running: Fallback to the alternative DNS servers happens globally. If the preferred DNS server fails to answer, all subsequent communication is with the alternative DNS servers.
    - Without the DNS Client service running: Any fallback to the alternative DNS servers happen locally, within each individual process that is making DNS queries. Different processes may be in different states, some talking to the preferred DNS server and some talking to alternative DNS servers.

[edit] Differences from other systems

Linux distributions and various versions of Unix have a generalized name resolver layer. The resolver can be controlled to use a **hosts** file or Network Information Service (NIS), by configuring the Name Service Switch.

[edit] Dynamic DNS Update client

Whilst DNS lookups read DNS data, DNS updates *write* them. Both workstations and servers running Windows attempt to send Dynamic DNS update requests to DNS servers.

Workstations running Windows attempt to register their names and their IP addresses with DNS servers, so that other machines may locate them by name. Prior to Windows Vista (and Windows Server 2008) this registration is performed by the *DHCP Client* service. It is thus necessary to run the DHCP Client service on pre-Vista machines, even if DHCP isn't being used to configure the machine in order to dynamically register a machine's name and address for DNS lookup. The DHCP Client service registers name and address data whenever they are changed (either manually by an administrator or automatically by the granting or revocation of a DHCP lease). In Microsoft Vista (and Windows Server 2008) Microsoft moved the registration functionallity from the *DHCP Client* service to the *DNS Client* service.

Servers running Microsoft Windows also attempt to register other information, in addition to their names and IP addresses, such as the locations of the LDAP and Kerberos services that they provide.

[edit] DNS server

Microsoft Windows server operating systems can run the *DNS Server* service. This is a monolithic DNS server that provides many types of DNS service, including caching, Dynamic DNS update, zone transfer, and DNS notification. DNS notification implements a push mechanism for notifying a select set of secondary servers for a zone when it is updated.

Microsoft's "DNS Server" service was first introduced in Windows NT 3.51 as an add-on with Microsoft's collection of BackOffice services (at the time was marked to be used for testing proposes only). Its code is a fork of ISC's BIND, version 4.3, and remains largely compatible with it including the format of all master files.

As of 2004, it was the fourth most popular DNS server (counting BIND version 9 separately from versions 8 and 4) for the publication of DNS data.[1]

Like various other DNS servers, Microsoft's DNS server supports different database *back ends*. Microsoft's DNS server supports two such back ends. DNS data can be stored either in *master files* (also known as *zone files*) or in the Active Directory database itself. In the latter case, since Active Directory (rather than the DNS server) handles the actual replication of the database across multiple machines, the database can be modified on any server ("multiple-master replication"), and the addition or removal of a *zone* will be immediately propagated to all other DNS servers within the appropriate Active Directory "replication scope". (Contrast this with BIND, where when such changes are made, the list of *zones*, in the `/etc/named.conf` file, has to be explicitly updated on each individual server.)

Microsoft's DNS server can be administered using either a graphical user interface, the "DNS Management Console", or a command line interface, the **dnscmd** utility.

[edit] Common issues

Prior to Windows Server 2003 and Microsoft Windows 2000 Service Pack 3, the most common problem encountered with Microsoft's DNS server was cache pollution. Although Microsoft's DNS Server had a mechanism for properly dealing with cache pollution, the mechanism was turned off by default. [2]

A longstanding well-known issue is incompatibility with BIND configuration files, in particular, the lack of support for DNS wildcards. This can be partially attributed to the fact that Microsoft's DNS server is based on BIND 4.3, before BIND added the support for DNS wildcards. *Loose Wildcarding* can be enabled. To create a wildcard character record, the *Dnscmd* command-line tool can be used. Also the support of IPv6 is implemented using a different technique from that of BIND 9, further driving even more incompatibilities between the two products.

In 2004, a common problem involved the feature of the Windows Server 2003 version of Microsoft's DNS server to use EDNS0, which a large number of firewalls could not cope with. [3]

[edit] See also

- Comparison of DNS server software

[edit] References

1. ^ Moore, Don (2004). "DNS server survey". http://mydns.bboy.net./survey/. Retrieved 2005-01-06.
2. ^ How to prevent DNS cache pollution
3. ^ Microsoft (April 6, 2006). "An external DNS query may cause an error message in Windows Server 2003". Microsoft. http://support.microsoft.com/kb/828731/. Retrieved 2006-05-08.

[edit] External links

- Microsoft Domain Name System (DNS)
- Jonathan de Boyne Pollard (2004). "Your firewall is preventing you from using EDNS0.". *Frequently Given Answers*. http://homepage.ntlworld.com./jonathan.deboynepollard/FGA/dns-edns0-and-firewalls.html. — the problems with EDNS0 and firewalls, and how to fix them
- Domain Name System Client Behavior in Windows Vista
- What's new in DNS in Windows Server 2008

How to Disable Client-Side DNS Caching in Windows XP and Windows Server 2003

View products that this article applies to.

This article was previously published under Q318803

For a Microsoft Windows 2000 version of this article, see 245437 .

On This Page

Expand all | Collapse all

SUMMARY

Windows contains a client-side Domain Name System (DNS) cache. The client-side DNS caching feature may generate a false impression that DNS "round robin" is not occurring from the DNS server to the Windows client computer. When you use the **ping** command to search for the same A-record domain name, the client may use the same IP address. This behavior is different from Microsoft operating systems earlier than Windows 2000. These operating systems do not include the client-side DNS caching feature. This article describes how to disable DNS caching.

Note This article refers to the client portion of DNS. Do not use this information for making changes to DNS servers.

DNS Client

From Process and Service wiki

Draft page (unreviewed)
Jump to: navigation, search

Contents

[hide]

General Information

If you attempt to "Diagnose" your network connection and a dialog box complains that the "DNS resolver failed to flush the cache," this service is the reason.

Only in extreme situations should you disable this service as caching DNS lookups reduces network traffic and makes internet surfing performance faster.

Windows 7

Default Description

The DNS Client service (dnscache) caches Domain Name System (DNS) names and registers the full computer name for this computer. If the service is stopped, DNS names will continue to be resolved. However, the results of DNS name queries will not be cached and the computer's name will not be registered. If the service is disabled, any services that explicitly depend on it will fail to start.

Additional Information

None at this time.

Default Startup Type

OS	SP0
Windows 7 Starter	Automatic (Started)
Windows 7 Home Basic	Automatic (Started)
Windows 7 Home Premium	Automatic (Started)
Windows 7 Professional	Automatic (Started)
Windows 7 Ultimate	Automatic (Started)
Windows 7 Enterprise	Automatic (Started)

Service Names

Service Name (registry): Dnscache

Display Name: DNS Client

Default Path and Command Line Options

C:\Windows\system32\svchost.exe -k NetworkService

Log On As

Account: Network Service

Dependencies

What service DNS Client needs to function properly:

- NetIO Legacy TDI Support Driver (S, HB, HP, P, U, E)
    - TCP/IP Protocol Driver (S, HB, HP, P, U, E)
- Network Store Interface Service (S, HB, HP, P, U, E)
    - NSI proxy service driver (S, HB, HP, P, U, E)

What other service require DNS Client to function properly:

- None (S, HB, HP, P, U, E)

Windows Vista

Default Description

The DNS Client service (dnscache) caches Domain Name System (DNS) names and registers the full computer name for this computer. If the service is stopped, DNS names will continue to be resolved. However, the results of DNS name queries will not be cached and the computer's name will not be registered. If the service is disabled, any services that explicitly depend on it will fail to start.

Additional Information

If you attempt to "repair" your network connection and a dialog box complains that the "DNS resolver failed to flush the cache," this service is the reason.

Only in extreme situations should you disable this service as caching DNS lookups reduces network traffic and makes internet surfing performance faster.

Default Startup Type

OS	SP0	SP1	SP2
Vista Home Basic	Automatic (Started)	Automatic (Started)	Automatic (Started)
Vista Home Premium	Automatic (Started)	Automatic (Started)	Automatic (Started)
Vista Business	Automatic (Started)	Automatic (Started)	Automatic (Started)
Vista Ultimate	Automatic (Started)	Automatic (Started)	Automatic (Started)
Vista Enterprise	Automatic (Started)	Automatic (Started)	Automatic (Started)

Service Names

Service Name (registry): Dnscache

Display Name: DNS Client

Default Path and Command Line Options

C:\Windows\system32\svchost.exe -k NetworkService

Log On As

Account: Network Service

Dependencies

What service DNS Client needs to function properly:

- NetIO Legacy TDI Support Driver (HB, HP, B, U)
  - TCP/IP Protocol Driver (HB, HP, B, U)

What other service require DNS Client to function properly:

- None (HB, HP, B, U)

Windows XP Pro x64

Default Description

Resolves and caches Domain Name System (DNS) names for this computer. If this service is stopped, this computer will not be able to resolve DNS names and locate Active Directory domain controllers. If this service is disabled, any services that explicitly depend on it will fail to start.

Additional Information

This service is not required for DNS lookups, but if it makes you happy to have it running, you may. If you attempt to "repair" your network connection and a dialog box complains that the "DNS resolver failed to flush the cache," this service is the reason. It is also needed if using IPSEC Service.

Default Startup Type

OS	SP0	SP1	SP2
XP Pro x64	?	?	Automatic

Service Names

Service Name (registry): DNScache

Display Name: DNS Client

Default Path and Command Line Options

C:\WINDOWS\system32\svchost.exe -k NetworkService

Log On As

Account: NT AUTHORITY\LocalService

Dependencies

What service DNS Client needs to function properly:

- TCP/IP Protocol Driver
    - IPSEC driver

What other service require DNS Client to function properly:

- None

Windows XP

Default Description

Resolves and caches Domain Name System (DNS) names for this computer. If this service is stopped, this computer will not be able to resolve DNS names and locate Active Directory domain controllers. If this service is disabled, any services that explicitly depend on it will fail to start.

Additional Information

This service is not required for DNS lookups, but if it makes you happy to have it running, you may. If you attempt to "repair" your network connection and a dialog box complains that the "DNS resolver failed to flush the cache," this service is the reason. It is also needed if using IPSEC Services.

Default Startup Type

OS	SP0	SP1	SP2	SP3
XP Home	Automatic (Started)	Automatic (Started)	Automatic (Started)	Automatic (Started)
XP MCE 2005	Automatic (Started)	Automatic (Started)	Automatic (Started)	Automatic (Started)
XP Pro	Automatic (Started)	Automatic (Started)	Automatic (Started)	Automatic (Started)
XP Tablet PC 2005	Automatic (Started)	Automatic (Started)	Automatic (Started)	Automatic (Started)

Service Names

SP3: Service Name (registry): Dnscache

Display Name: DNS Client

Default Path and Command Line Options

C:\WINDOWS\system32\svchost.exe -k NetworkService

Log On As

Account: NT AUTHORITY\LocalService

SP3: Account: NT AUTHORITY\NetworkService

Dependencies

What service DNS Client needs to function properly:

- TCP/IP Protocol Driver (H, M, P, T)
    - IPSEC driver (H, M, P, T)

What other service require DNS Client to function properly:

- None (H, M, P, T)

Additional Reading

- DNS: http://en.wikipedia.org/wiki/Domain_Name_System
- Cache: http://en.wikipedia.org/wiki/Cache

Retrieved from "http://wiki.blackviper.com/wiki/DNS_Client"

Troubleshooting DNS servers

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Troubleshooting DNS servers

What problem are you having?

- The DNS server is not responding to clients.
- The DNS server does not resolve names correctly.
- The DNS server appears to be affected by a problem for reasons not described above.

The DNS server is not responding to clients.
Cause: The DNS server is affected by a network failure.

Solution: Verify that the server computer has a valid functioning network connection. First, check that related client hardware (cables and network adapters) are working properly at the client using basic network and hardware troubleshooting steps.

If the server hardware appears to be prepared and functioning properly, check that it has network connectivity by pinging other computers or routers (such as its default gateway) that are used and available on the same network as the affected DNS servers.

See also: Test a TCP/IP configuration by using the ping command.

Cause: The DNS server is reachable through basic network testing but is not responding to DNS queries from clients.

Solution: If the DNS client can ping the DNS server computer, verify that the DNS server is started and able to listen to and respond to client requests. Try using the **nslookup** command to test whether the server can respond to DNS clients.

See also: Verify DNS server responsiveness using the nslookup command; Start or stop a DNS server.

Cause: The DNS server has been configured to limit service to a specific list of its configured IP addresses. The IP address originally used in testing its responsiveness is not included in this list.

Solution: If the server was previously configured to restrict the IP addresses for which it responds to queries, it is possible that the IP address being used by clients to contact it is not in the list of restricted IP addresses permitted to provide service to clients.

Try testing the server for a response again, but specify a different IP address known to be in the restricted interfaces list for the server. If the DNS server responds for that address, add the missing server IP address to the list.

See also: Verify DNS server responsiveness using the nslookup command; Restrict a DNS server to listen only on selected addresses.

Cause: The DNS server has been configured to disable the use of its automatically created default reverse lookup zones.

Solution: Verify that automatically created reverse lookup zones have been created for the server or that advanced configuration changes have not been previously made to the server.

By default, DNS servers automatically create the following three standard reverse lookup zones based on Request for Comments (RFC) recommendations:

These zones are created with common IP addresses covered by these zones that are not useful in a reverse lookup search (0.0.0.0, 127.0.0.1, and 255.255.255.255). By being authoritative for the zones corresponding to these addresses, the DNS service avoids unnecessary recursion to root servers in order to perform reverse lookups on these types of IP addresses.

It is possible, although unlikely, that these automatic zones are not created. This is because disabling the creation of these zones involves advanced manual configuration of the server registry by a user.

To verify that these zones have been created, do the following:

1. Open the DNS console.

2. From the **View** menu, click **Advanced**.

3. In the console tree, click **Reverse Lookup Zones**.

   Where?

   - **DNS**/*applicable DNS server*/**Reverse Lookup Zones**

4. In the details pane, verify that the following reverse lookup zones are present:

   - 0.in-addr.arpa
   - 127.in-addr.arpa
   - 255.in-addr.arpa

See also: Open the DNS console; DNS RFCs.

Cause: The DNS server is configured to use a non-default service port, such as in an advanced security or firewall configuration.

Solution: Verify that the DNS server is not using a non-standard configuration.

This is a rare but possible cause. By default, the **nslookup** command sends queries to targeted DNS servers using User Datagram Protocol (UDP) port 53. If the DNS server is located on another network only reachable through an intermediate host (such as a packet-filtering router or proxy server), the DNS server might use a non-standard port to listen for and receive client requests.

If this situation applies, determine whether any intermediate firewall or proxy server configuration is intentionally used to block traffic on well-known service ports used for DNS. If not, you might be able to add such a packet filter onto these configurations to permit traffic to standard DNS ports.

Also, check the DNS server event log to see if Event ID 414 or other critical service-related events have occurred which might indicate why the DNS server is not responding.

See also: DNS server log reference; View the DNS server system event log; Microsoft Windows Deployment and Resource Kits.

The DNS server does not resolve names correctly.
Cause: The DNS server provides incorrect data for queries it successfully answers.

Solution: Determine the cause of the incorrect data for the DNS server.

Some of the most likely causes include the following:

- Resource records (RRs) were not dynamically updated in a zone.
- An error was made when manually adding or modifying static resource records in the zone.
- Stale resource records in the DNS server database, left from cached lookups or zone records not updated with current information or removed when they are no longer needed.

To help prevent the most common types of problems, be sure to first review best practices for tips and suggestions on deploying and managing your DNS servers. Also, follow and use the checklists appropriate for installing and configuring DNS servers and clients based on your deployment needs.

If you are deploying DNS for Active Directory, note new directory integration features. These features can cause some differences for DNS server defaults when the DNS database is directory-integrated, that differ from those used with traditional file-based storage.

Many DNS server problems start with failed queries at a client, so it is often good to start there and troubleshoot the DNS client first.

See also: DNS best practices; DNS Checklists; Troubleshooting DNS clients; Modify an existing resource record in a zone; Clear the server names cache; Modifying server defaults.

Cause: The DNS server does not resolve names for computers or services outside of your immediate network, such as those located on external networks or the Internet.

Solution: The server has a problem based on its ability to correctly perform recursion. Recursion is used in most DNS configurations to resolve names that are not located within the configured DNS domain name used by the DNS servers and clients.

If a DNS server fails to resolve a name for which it is not authoritative, the cause is usually a failed recursive query. Recursive queries are used frequently by DNS servers to resolve remote names delegated to other DNS zones and servers.

For recursion to work successfully, all DNS servers used in the path of a recursive query must be able to respond to and forward correct data. If not, a recursive query can fail for any of the following reasons:

- The recursive query times out before it can be completed.
- A remote DNS server fails to respond.

- A remote DNS server provides incorrect data.

If a server fails a recursive query for a remote name, review the following possible causes to troubleshoot the problem. If you do not understand recursion or the DNS query process, review conceptual topics in Help to better understand the issues involved.

See also: How DNS query works.

Cause: The DNS server is not configured to use other DNS servers to assist it in resolving queries.

Solution: Check whether the DNS server can use both forwarders and recursion.

By default, all DNS servers are enabled to use recursion, although the option to disable its use is configurable using the DNS console to modify advanced server options. The other possibility where recursion might be disabled is if the server is configured to use forwarders and recursion has been specifically disabled for that configuration.

Note

- If you disable recursion on the DNS server, you will not be able to use forwarders on the same server.

See also: Disable recursion on the DNS server; Configure a DNS server to use forwarders.

Cause: Current root hints for the DNS server are not valid.

Solution: Check whether server root hints are valid.

If configured and used correctly, root hints always should point to DNS servers authoritative for the zone containing the domain root and top-level domains.

By default, DNS servers are configured to use root hints appropriate to your deployment, based on the following available choices when using the DNS console to configure a server:

1. If the DNS server is installed as the first DNS server for your network, it is configured as a root server.

   For this configuration, root hints are disabled at the server because the server is authoritative for the root zone.

2. If the installed server is an additional DNS server for your network, you can direct the Configure DNS Server Wizard to update its root hints from an existing DNS server on the network.

3. If you do not have other DNS servers on your network but still need to resolve Internet DNS names, you can use the default root hints file which includes a list of Internet root servers authoritative for the Internet DNS namespace.

See also: Update root hints on the DNS server; Updating root hints.

Cause: The DNS server does not have network connectivity to the root servers.

Solution: Test for connectivity to the root servers.

If root hints appear to be configured correctly, verify that the DNS server used in a failed query can ping its root servers by IP address.

If a ping attempt to one root server fails, it might indicate that an IP address for that root server has changed. Reconfiguration of root servers, however, is uncommon.

A more likely cause is a full loss of network connectivity or in some cases, poor network performance on the intermediate network links between the DNS server and its configured root servers. Follow basic TCP/IP network troubleshooting steps to diagnose connections and determine whether this is the problem.

By default, the DNS service uses a recursive time-out of 15 seconds before failing a recursive query. Under normal network conditions, this time-out does not need to be changed. If performance warrants it, however, you can increase this value.

To review additional performance related information on DNS queries, you can enable and use the DNS server debug log file, Dns.log, which can provide extensive information about some types of service-related events.

See also: Test a TCP/IP configuration by using the ping command; Using server debug logging options; View a DNS server debug log file; Tuning advanced server parameters.

Cause: Other problems exist with updating DNS server data, such as an issue related to zones or dynamic updates.

Solution: Determine whether the problem is related to zones. As needed, Troubleshoot any issues in this area, such as possible failure of zone transfer.

See also: Troubleshooting dynamic updates; Troubleshooting zone problems.

The DNS server appears to be affected by a problem for reasons not described above.
Cause: My problem is not described above.

Solution: Search TechNet at the Microsoft Web site for the latest technical information that could relate to the problem. If necessary, you can obtain information and instructions that pertain to your problem or issue.

If you are connected to the Internet, the latest operating system updates are available at the Microsoft Web site.

To obtain the latest service pack updates for Windows NT Server, see the Microsoft Web site.

See also: DNS updated technical information; DNS; Using the Windows Deployment and Resource Kits.

Tags : Add a tag

Community Content

Add new content Annotations

Troubleshooting DNS Problems with Preparation Wizard

Tool to Troubleshoot Network Problems ... mwatson677 | Edit false | Show History

Use a great free Microsoft tool developed by the Microsoft Essential Business Server team (http://blogs.technet.com/essentialbusinessserver/). The tool scans your network, identifies various networking problems (DNS, AD replication, AD configuration, configuration of network adapters, etc.) and provides links to knowledge based articles that explain how to correct these issues.

Microsoft built this tool for customers who are considering deploying Essential Business Server 2008 (http://www.microsoft.com/ebs/en/us/overview.aspx) to prepare their environment for Essential Business Server (hence the name: Preparation Wizard).

But, this tool can be used by anyone with Active Directory in their network who would like to verify the health of their environment. The tool runs over 100 different checks which are based on most common issues resolved by Microsoft Customer Support Services over the past 10 years!

The tool is specifically designed for mid-sized networks (**25-300 PCs**), and does not change any settings in your network, so it is safe to run at any time. Unlike many other known tools which simply dump large amounts of networking data collected from a single source (such as event logs, for instance), this tool is able to gather data from many different areas(Active Directory, DNS, SYSVOL, event logs, etc.), cross reference that data, and make conclusions about the overall health of the network.

Try it today. Go get it – it's FREE!
http://www.microsoft.com/ebs/en/us/preparation.aspx

Tags : contentbug (x) problems (x) troubleshooting (x) wizard (x) Add a tag

DNS Error, Event ID 4015 Romelo5179 ... Thomas Lee | Edit false | Show History

I have 2 domain controllers. The second was setup as a backup with DNS installed on it. Every 48-72 hours, I get the Event ID 4015 AD error message and then continuous errors until I restart the server. As soon as I restart the server, the problem is resolved and it looks like nothing is wrong. But then, 48-72 hours later, Event ID 4015 AD error again and it does the same thing as before until I restart. What could be causing this? I have read so many tech notes and everything I can find on this error and still the same preoblem occurs. Any ideas?

I am using Server 2003.

[tfl - 04 05 09] You should post questions like this to the Technet Forums at http://forums.microsoft.com/technet or the MS Newsgroups at http://www.microsoft.com/communities/newsgroups/en-us/. You are much more likely get a quick response using the forums than through the Community Content. For specific help about:
Exchange :
http://groups.google.com/groups/dir?sel=usenet%3Dmicrosoft.public.exchange%2C&;;
SQL Server :
http://groups.google.com/groups/dir?sel=usenet%3Dmicrosoft.public.sqlserver%2C&;;
Windows :
http://groups.google.com/groups/dir?sel=usenet%3Dmicrosoft.public.windows%2C&;;
Windows Server :
http://groups.google.com/groups/dir?sel=usenet%3Dmicrosoft.public.windows.server%2C&;;
Virtual Server :
http://groups.google.com/group/microsoft.public.virtualserver/topics?lnk
Full Public :
http://groups.google.com/groups/dir?sel=usenet%3Dmicrosoft.public%2C&;;

Tags : needsforum (x) needsnewsgroup (x) Add a tag

Flag as ContentBug

Troubleshooting DNS Problems with Preparation Wizard Tool to Troubleshoot Network Problems | Edit false | Show History

http://www.microsoft.com/ebs/en/us/preparation.aspx

Tags : dns (x) network (x) preparation (x) problems (x) server (x) troubleshooting (x) wizard (x) Add a tag

Flag as ContentBug

Troubleshooting DNS Problems with Preparation

http://technet.microsoft.com/en-us/library/cc787191(WS.10).aspx

Manually create DNS entry

Operating Systems - Windows

Manually create a DNS entry for Server 2003

Sometimes you will want to know for certain that a DNS server is resolving a hostname correctly. Making a manual entry might be the only way!

Create a DNS entry

This article assumes that DNS is installed on your server.

1. Log on to your Server 2003.
2. Go to **Start > Administrative Tools > DNS**.
3. Browse to **YourServerName > Forward Lookup Zones > YourDomain**.
4. Right-click on **YourDomain** and select **New Host (A)...**
5. Enter the **Name** and **IP address** of the host that you want to add.
6. Leave the **Create associated pointer (PTR) record** checked.
7. Click **Add Host**.
8. *Click **OK** if you get a warning, and click **Done**.*

Back

More articles of interest...

- Configure Ubuntu Servers
- Configure Server 2003 Enterprise
- File Server

http://technet.microsoft.com/en-us/library/bb727020.aspx

Managing DNS Server Configuration and Security

from Chapter 19, Microsoft Windows 2000 Administrator's Pocket Consultant by William R. Stanek.

You use the Server Properties dialog box to manage the general configuration of DNS servers. Through it, you can enable and disable IP addresses for the server and control access to DNS servers outside the organization. You can also configure monitoring, logging, and advanced options.

Enabling and Disabling IP Addresses for a DNS Server

By default, multihomed DNS servers respond to DNS requests on all available network adapters and the IP addresses they're configured to use.

Through the DNS console, you can specify that the server can only answer requests on specific IP addresses. To do this, follow these steps:

1. In the DNS console, right-click the server you want to configure and then from the pop-up menu, choose Properties.
2. In the Interfaces tab, shown in Figure 19-14, select Only The Following IP Addresses and then type the IP addresses that should respond to DNS requests. Only these IP addresses will be used for DNS. All other IP addresses on the server will be disabled for DNS.

Controlling Access to DNS Servers Outside the Organization

Restricting access to zone information allows you to specify which internal and external servers can access the primary server. For external servers, this controls which servers can get in from the outside world. You can also control which DNS servers within your organization can access servers outside it. To do this, you need to set up DNS forwarding within the domain.

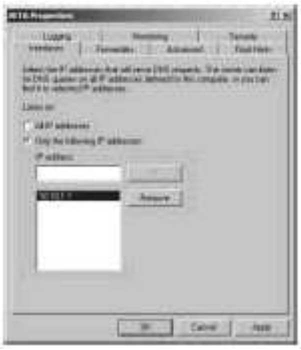

Figure 19-14: Use the Interfaces tab to set the IP addresses that should handle DNS requests and responses.

With DNS forwarding, you configure DNS servers within the domain as

- **Nonforwarders** Servers that must pass DNS queries they can't resolve on to designated forwarding servers. These servers essentially act like DNS clients to their forwarding servers.
- **Forwarding-only** Servers that can only cache responses and pass requests on to forwarders. This is also known as a *caching-only* DNS server.

- **Forwarders** Servers that receive requests from nonforwarders and forwarding-only servers. Forwarders use normal DNS communication methods to resolve queries and to send responses back to other DNS servers.

Note: The root server for a domain can't be configured for forwarding. But all other servers can be configured for forwarding.

Creating Nonforwarding DNS Servers

To create a nonforwarding DNS server, follow these steps:

1. In the DNS console, right-click the server you want to configure and then from the pop-up menu, choose Properties.
2. In the Forwarders tab, select Enable Forwarders.
3. Enter the IP addresses of the network's forwarders.
4. Set the Forward Time Out. This value controls how long the server tries to query the server if it gets no response. When the Forward Time Out interval passes, the server tries the next forwarder on the list. The default is 0 seconds. Click OK.

Creating Forwarding-Only Servers

To create a forwarding-only server, follow these steps:

1. In the DNS console, right-click the server you want to configure and then from the pop-up menu, choose Properties.
2. In the Forwarders tab, select Enable Forwarders and then select Operate As Slave Server.
3. Enter the IP addresses of the network's forwarders.
4. Set the Forward Time Out. This value controls how long the server tries to query the server if it gets no response. When the Forward Time Out interval passes, the server tries the next forwarder on the list. The default is 0 seconds. Click OK.

Creating Forwarders Servers

Any DNS server that isn't designated as a nonforwarder or a forwarding-only server will act as a forwarder. Thus, on the network's designated forwarders, you should make sure that Enable Forwarders and Operate As Slave Server are *not* selected.

Logging DNS Activity

You normally use the DNS Server event log to track DNS activity on a server. This log records all applicable DNS events and is accessible through the Event View node in Computer Management. If you're trying to troubleshoot DNS problems, it's sometimes useful to configure a temporary debug log to track certain types of DNS events. To do this, follow these steps:

1. In the DNS console, right-click the server you want to configure and then from the pop-up menu, choose Properties.
2. In the Logging tab, shown in Figure 19-15, select the events you want to track temporarily. These events are logged in %SystemRoot%\System32\Dns\ Dns.log by default.

3. Click OK. When you're finished debugging, turn off logging by clearing any of the selected check boxes in the Logging tab.

Monitoring DNS Server

Windows 2000 has built-in functionality for monitoring DNS server. You can configure monitoring to occur manually or automatically by completing the following steps:

1. In the DNS console, right-click the server you want to configure and then from the pop-up menu, choose Properties.
2. Select the Monitoring tab, shown in Figure 19-16. You can perform two types of tests. To test DNS resolution on the current server, select A Simple Query Against This DNS Server. To test DNS resolution in the domain, select A Recursive Query To Other DNS Servers.

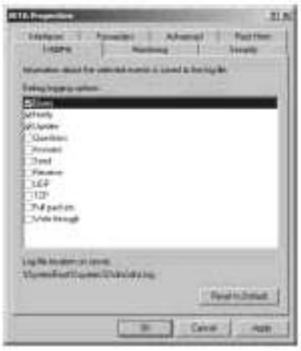

Figure 19-15: Select the events you want to log, and then click OK. Don't forget to clear these events after you've finished debugging.

3. You can perform a manual test by clicking Test Now or schedule the server for automatic monitoring by selecting Perform Automatic Testing At The Following Interval and then setting a time interval in seconds, minutes, or hours.

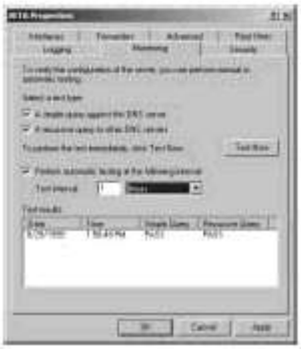

Figure 19-16: You can configure a DNS server for manual or automatic monitoring. Monitoring is useful to ensure that DNS resolution is configured properly.

Real World If you're actively troubleshooting a DNS problem, you may want to configure testing to occur every 10–15 seconds. This will provide a rapid succession of test results. If you're monitoring DNS for problems as part of your daily administrative duties, you'll want a longer time interval, such as two or three hours.

4. The results of testing are shown in the Test Results area. You'll see a date and time stamp indicating when the test was performed and a result, such as Pass or Fail. While a single failure may be the result of a temporary outage, multiple failures normally indicate a DNS resolution problem.

Integrating WINS with DNS

You can integrate DNS with WINS. WINS integration allows the server to act as a WINS server or to forward WINS requests to specific WINS servers. When you configure WINS and DNS to work together, you can configure forward lookups using NetBIOS computer names, reverse lookups using NetBIOS computer names, caching and time-out values for WINS resolution, and full integration with NetBIOS scopes.

Configuring WINS Lookups in DNS

When you configure WINS lookups in DNS, the leftmost portion of the fully qualified domain name can be resolved using WINS. The procedure works like this: The DNS server looks for an address record for the fully qualified domain name. If a record is found, the server uses the record to resolve the name using only DNS. If a record isn't found, the server extracts the leftmost portion of the name and uses WINS to try to resolve the name (as a NetBIOS computer name). You configure WINS lookups in DNS by doing the following:

1. In the DNS console, right-click the domain you want to update and then from the pop-up menu, choose Properties.
2. Click the WINS tab, shown in Figure 19-17.
3. Select Use WINS Forward Lookup and then type the IP addresses of the network's WINS servers. You must specify at least one WINS server.
4. If you want to ensure that the WINS record on this server isn't replicated to other DNS servers in zone transfers, select Do Not Replicate This Record. Selecting this option is useful to prevent errors and transfer failures to non-Microsoft DNS servers. Click OK.

Configuring Reverse WINS Lookups in DNS

When you configure reverse WINS lookups in DNS, the IP address of the host can be resolved to a NetBIOS computer name. The procedure works like this: The DNS server looks for a pointer record for the specified IP address. If a record is found, the server uses the record to resolve the fully qualified domain name.

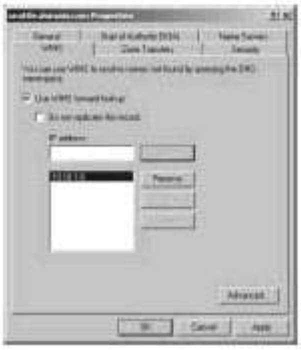

Figure 19-17: Use the WINS tab to configure WINS lookups in DNS.

If a record isn't found, the server sends a request to WINS, and, if possible, WINS returns the NetBIOS computer name for the IP address and the host domain is appended to this computer name.

You configure reverse WINS lookups in DNS by doing the following:

1. In the DNS console, right-click the subnet you want to update and then from the pop-up menu, choose Properties.

2. Click the WINS-R tab, shown in Figure 19-18.

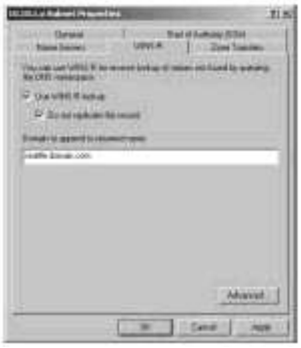

You can use WINS-R for reverse lookup of names not found by querying the DNS namespace.

☑ Use WINS-R lookup

☑ Do not replicate this record

Domain to append to returned name:

seattle.domain.com

Figure 19-18: Use the WINS-R tab to configure WINS reverse lookups in DNS.

3. Select Use WINS-R Lookup, and then, if you wish, select Do Not Replicate This Record. As with forward lookups, you usually don't want to replicate the WINS-R record to non-Microsoft DNS servers.
4. In the Domain To Append To Returned Name field, type the host domain information. The domain is appended to the computer name returned by WINS. For example, if you type **seattle.domain.com** and WINS returns the NetBIOS computer name gamma, the DNS server will combine the two values and return gamma.seattle.domain.com.
5. Click OK.

Setting Caching and Time-Out Values for WINS in DNS

When you integrate WINS and DNS, you should also set WINS caching and time-out values. The caching value determines how long records returned from WINS are valid. The time-out value determines how long DNS should wait for a response from WINS before timing out and returning an error. These values are set for both forward and reverse WINS lookups.

You set caching and time-out values for WINS in DNS by doing the following:

1. In the DNS console, right-click the domain or subnet you want to update and then from the pop-up menu, choose Properties.
2. Select the WINS or WINS-R tab, as appropriate, and then click Advanced. This opens the dialog box shown in Figure 19-19.
3. Set the caching and time-out values using the Cache Time-Out field and the Lookup Time-Out field. By default, DNS caches WINS records for 15 minutes and times out after 2 seconds. For most networks, you should increase these values. Sixty minutes for caching and three seconds for time-outs may be better choices.
4. Click OK. Repeat this process for other domains and subnets, as necessary.

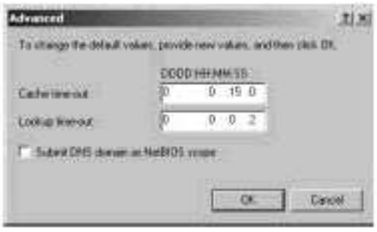

To change the default values, provide new values, and then click OK.

DDDD HH:MM:SS

Cache time-out: 0 0 15 0

Lookup time-out: 0 0 0 2

☐ Submit DNS domain as NetBIOS scope

Figure 19-19: In the Advanced dialog box, set caching and time-out values for DNS.

Configuring Full Integration with NetBIOS Scopes

When you configure full integration, lookups can be resolved using NetBIOS computer names and NetBIOS scopes. Here, a forward lookup works like this: The DNS server looks for an address record for the fully qualified domain name. If it finds a record, the server uses the record to resolve the name using only DNS. If it doesn't find a record, the server extracts the leftmost portion of the name as the NetBIOS computer name and the remainder of the name as the NetBIOS scope. These values are then passed to WINS for resolution.

You configure full integration of WINS and DNS by doing the following:

1. In the DNS console, right-click the domain or subnet you want to update, and then from the pop-up menu, choose Properties.
2. Select the WINS or WINS-R tab, as appropriate, and then click Advanced.
3. In the Advanced dialog box, select Submit DNS Domain As NetBIOS Scope.
4. Click OK. Repeat this process for other domains and subnets, as necessary.

Before you use this technique, make sure that the NetBIOS scope is properly configured on the network. You should also make sure that a consistent naming scheme is used for all network computers. Because NetBIOS is case-sensitive, queries resolve only if the case matches exactly. Note also that if the domain has subdomains, the subdomains must be delegated the authority for name services in order for WINS and DNS integration to work properly.

from Microsoft Windows 2000 Administrator's Pocket Consultant by William R. Stanek. Copyright © 1999 Microsoft Corporation.

Dynamic Host Configuration Protocol (DHCP)

DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

DHCP

Dynamic Host Configuration Protocol (DHCP) is an IP standard designed to reduce the complexity of administering address configurations by using a server computer to centrally manage IP addresses and other related configuration details used on your network. The Microsoft® Windows Server 2003 family provides the DHCP service, which enables the server computer to perform as a DHCP server and configure DHCP-enabled client computers on your network as described in the current DHCP draft standard, RFC 2131.

DHCP includes Multicast Address Dynamic Client Assignment Protocol (MADCAP) which is used to perform multicast address allocation. When registered clients are dynamically assigned IP addresses through MADCAP, they can participate efficiently in the data stream process, such as for real-time video or audio network transmissions.

- Before installing a DHCP or MADCAP server, see DHCP Checklists.
- To find features that have been moved in the Windows Server 2003 family, see New ways to do DHCP tasks.
- For tips about using DHCP, see DHCP Best Practices.
- For help with specific tasks, see DHCP How To....
- For general background information, see DHCP Concepts.
- For problem-solving instructions, see DHCP Troubleshooting.

How to install and configure a DHCP server in an Active Directory domain in Windows 2000

View products that this article applies to.

System Tip This article applies to a different version of Windows than the one you are using. ✕

Content in this article may not be relevant to you.Visit the Windows XP Solution Center

This article was previously published under Q300429

Notice

This article applies to Windows 2000. **Support for Windows 2000 ends on July 13, 2010.** The Windows 2000 End-of-Support Solution Centeris a starting point for planning your migration strategy from Windows 2000. For more information see the Microsoft Support Lifecycle Policy.

Notice

This article applies to Windows 2000. **Support for Windows 2000 ends on July 13, 2010.** The Windows 2000 End-of-Support Solution Centeris a starting point for planning your migration strategy from Windows 2000. For more information see the Microsoft Support Lifecycle Policy.

On This Page

Expand all | Collapse all

SUMMARY

This step-by-step article describes how to build and configure a new Windows 2000 DHCP Server in a Windows 2000 Active Directory domain. The Windows 2000 DHCP service provides clients with IP addresses, and information such as the location of their default gateway, DNS servers, and WINS servers.

↑Back to the top

Installing the DHCP Service

You can install DHCP either during or after the initial installation of Windows 2000 Server or Advanced Server, although there must be a working DNS in the environment. To validate your DNS server, click **Start**, click **Run**, type **cmd**, press ENTER, type **ping *friendly name of an existing DNS server in your environment***, and then press ENTER. An unsuccessful reply generates an "Unknown Host My *DNS server name*" message.

To install the DHCP Service on an existing Windows 2000 Server:

1. Click **Start**, click **Settings**, and then click **Control Panel**.
2. Double-click **Add/Remove Programs**, and then click **Add/Remove Windows Components**.
3. In the **Windows Component Wizard**, click **Networking Services** in the **Components** box, and then click **Details**.
4. Click to select the **Dynamic Host Configuration Protocol (DHCP)** check box if it is not already selected, and then click **OK**.
5. In the **Windows Components Wizard**, click **Next** to start Windows 2000 Setup. Insert the Windows 2000 Advanced Server CD-ROM into the CD-ROM drive if you are prompted to do so. Setup copies the DHCP server and tool files to your computer.
6. When Setup is complete, click **Finish**.

↑Back to the top

Configuring the DHCP Service

After you install and start the DHCP service, you must create a scope (a range of valid IP addresses that are available for lease to the DHCP clients). Each DHCP server in your environment should have at least one scope that does not overlap with any other DHCP server scope in your environment. In Windows 2000, DHCP servers within an Active Directory domain environment must be authorized to prevent rogue DHCP servers from coming online and authorizing a DHCP Server.

When you install and configure the DHCP service on a domain controller, the server is typically authorized the first time that you add the server to the DHCP console. However, when you install and configure the DHCP service on a member server, you need to authorize the DHCP server.

Note A stand-alone DHCP server cannot be authorized against an existing Windows Active Directory.

To authorize a DHCP server:

1. Click **Start**, click **Programs**, click **Administrative Tools**, and then click **DHCP**.

   Note You must be logged on to the server with an account that is a member of the Enterprise Administrators group.
2. In the console tree of the DHCP snap-in, select the new DHCP server. If there is a red arrow in the bottom-right corner of the server object, the server has not yet been authorized.
3. Right-click the server, and then click **Authorize**.
4. After a few moments, right-click the server again and then click **Refresh**. The server should display a green arrow in the bottom-right corner to indicate that the server has been authorized.

To create a new scope:

1. Click **Start**, click **Programs**, point to **Administrative Tools**, and then click **DHCP**.

   Note In the console tree, select the DHCP server on which you want to create the new DHCP scope.
2. Right-click the server, and then click **New Scope**. In the New Scope Wizard, click **Next**, and then type a name and description for the scope. This can be any name that you choose, but it should be descriptive enough to identify the purpose of the scope on your network. For example, you might use Administration Building Client Addresses.
3. Type the range of addresses that can be leased as part of this scope, for example, a starting IP address of 192.168.100.1 to an ending address of 192.168.100.100. Because these addresses are given to clients, they should all be valid addresses for your network and not currently in use. If you want to use a different subnet mask, type the new subnet mask. Click **Next**.
4. Type any IP addresses that you want to exclude from the range you entered. This includes any addresses that may have already been statically assigned to various computers in your organization. Click **Next**.
5. Type the number of days, hours, and minutes before an IP address lease from this scope expires. This determines the length of time that a client can hold a leased address without

renewing it. Click **Next** to select **Yes, I want to configure these options now**, and then extend the wizard to include settings for the most common DHCP options. Click **Next**.

6. Type the IP address for the default gateway that should be used by clients that obtain an IP address from this scope. Click **Add** to place the default gateway address into the list, and then click **Next**.

   Note When DNS servers already exist on your network, type your organization's domain name in **Parent domain**. Type the name of your DNS server, and then click **Resolve** to ensure that your DHCP server can contact the DNS server and determine its address. Then click **Add** to include that server in the list of DNS servers that are assigned to the DHCP clients. Click **Next**.
7. Click **Yes, I want to activate this scope now**, to activate the scope and allow clients to obtain leases from it, and then click **Next**. Click **Finish**.

↑Back to the top

Troubleshooting

- **Clients are unable to obtain an IP address**
  If a DHCP client does not have a configured IP address, it generally means that the client has not been able to contact a DHCP server. This is either because of a network problem or because the DHCP server is unavailable. If the DHCP server has started and other clients have been able to obtain a valid address, verify that the client has a valid network connection and that all related client hardware devices (including cables and network adapters) are working properly.
- **The DHCP server is unavailable**
  When a DHCP server does not provide leased addresses to clients, it is often because the DHCP service has failed to start. If this is the case, the server may not have been authorized to operate on the network. If you were previously able to start the DHCP service, but it has since stopped, use Event Viewer to check the system log for any entries that may explain the cause.

  Note To restart the DHCP service, click **Start**, click **Run**, type **cmd**, and then press ENTER. Type **net start dhcpserver**, and then press ENTER.

↑Back to the top

Using multicast scopes

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Using multicast scopes

Multicast scopes are supported through the use of Multicast Address Dynamic Client Allocation Protocol (MADCAP), a proposed standard protocol for performing multicast address allocation. The MADCAP protocol describes how multicast address allocation (or MADCAP) servers can dynamically provide IP addresses to other computers (MADCAP clients) on your network.

Typically, a MADCAP client might also be a multicast server (MCS) used to support IP multicasting. An MCS manages the shared or group use of the allocated multicast IP address and streams data traffic to members that share the use of the specified group address.

Once an MCS is configured and allocated a group IP address to use, any multicast clients that have registered their membership with the MCS can receive streams sent to this address. By registering with the MCS, clients can participate efficiently in the stream process, such as for real-time video or audio network transmissions. The MCS also manages the multicast group list, updating its membership and status so that multicast traffic is received by all current members.

Background on multicast addresses

Ordinarily, you use DHCP scopes to provide client configurations by allocating ranges of IP addresses from the standard address classes, Class A, B, or C. By using DHCP scopes you can assign IP addresses from the ranges provided by these classes for your DHCP clients to be configured to use unicast (or point-to-point) directed communication between other TCP/IP networked computers.

The multicast address range uses an additional address class, Class D, that includes IP addresses that range from 224.0.0.0 to 239.255.255.255 for use in IP multicasting. Addresses in this class are used for multicasting only and not for regular DHCP scopes.

In all TCP/IP networks, each host is required to first be configured with its own IP address, taken from one of the standard address classes. You must assign this required unicast IP address before you can configure a host to support and use secondary IP addresses, such as a multicast IP address.

Another difference between unicast and multicast addresses, is that a group of TCP/IP host computers are intended to be able to share the use of a multicast IP address. This is not normally the case with unicast IP addresses, which are intended to be assigned individually to only one configured host and not shared with other hosts.

When the destination address for an IP datagram is an IP multicast address, the datagram is forwarded to all members of a *multicast group*, which is a set of zero or more hosts identified by the address. The membership of a multicast group is dynamic in that individual hosts can join or leave the group at any time.

Membership and use of multicast groups is unrestricted and can be compared to membership and use of a group e-mail address: Group membership can be any size, and hosts can be members of many multicast groups.

You can permanently reserve multicast group addresses or temporarily assign and use them as needed on your network. For a permanent group IP address to be reserved for Internet use, you must register it with the Internet Assigned Numbers Authority (IANA).

For multicast IP addresses not permanently reserved with the IANA, all Class D addresses that remain unreserved can then be used dynamically to assign and form temporary multicast groups. These temporary groups can exist as long as one or more hosts on the network are configured with the group address and actively share its use.

For more information, see Understanding Multicasting or Internet Group Management Protocol (IGMP).

Determining the ranges to use for multicast scopes

When deciding the IP address ranges to use for multicast scopes on your MADCAP server, there are two overall best practices recommended by the Multicast Allocation (MALLOC) working group, an IETF team of industry volunteers who help establish multicast address allocation standards. These practices include:

- **Administrative scoping** This is most useful when you are using multicast IP addresses privately in your own network. It is similar to private IP addressing, as used currently in the unicast IP address spaces (such as the use of the 10.0.0.0 network address space). It is fully discussed in RFC 2365, "Administratively Scoped IP Multicast."

  With administratively scoped multicast IP, the range most recommended that you begin with is the 239.192.0.0 range. This range is known as the IPv4 Organization Local Scope and has a subnet mask of 255.252.0.0 (14-bits in length). It is intended for use by an entire organization setting multicast scopes privately for its own internal or organizational use. Starting with this address, you can create a considerable number of addresses--up to 2^{18}, or 262,144 group addresses--for use in all subnets within your organization's network.

  For more information, refer directly to the RFC, which can be obtained at the Request for Comments Web site.

- **Global scoping** This practice is most useful when you are using multicast group IP addresses in a public network address space, particularly the Internet.

  Because most organizations requiring one or more public addresses likely have already been assigned some public unicast IP addresses, a subnet allocation scheme has been proposed. For MADCAP, the 233.0.0.0 range of the Class D address space is recommended for use as a global scope range.

  When the 233.0.0.0 range is used, the allocation of global addresses works in the following way:

  1. A network registry, such as the Internet Assigned Numbers Authority (IANA), allocates and reserves the first 8 bits of the range (for example, the "233" portion of this range).

  2. The next 16 bits (the two middle octet numbers of addresses in this range) are based on a previously assigned Autonomous System (AS) number. This number is recorded with the applicable Internet Assigned Numbers Authority (IANA) registry for your region.

     If you already register your IP addresses with a regional network registry, you might know your AS number. You can also look it up easily using the Whois database system on the Internet. If you are obtaining IP addresses through an Internet service provider (ISP), you might contact them to find out what your AS number is.

     AS numbers are allocated to the regional registries by the IANA. If you or your ISP are located in the United States, you can apply directly to the American Registry for Internet Numbers (ARIN) to obtain an AS number. For more information, see the ARIN Web site.

     For other regions, AS numbers can also be obtained from appropriate regional registries. Other regional registries include the following:

ARIN or <hostmaster@arin.net> for the Americas, Caribbean, and Africa.

RIPE-NCC <ncc@ripe.net> for Europe.

AP-NIC <admin@apnic.net> for the Asia-Pacific region.

1. The last 8 bits in the address are local use bits.

These provide the IP address range from which to configure any multicast scopes for group addresses you want to use publicly on the Internet. With this global scoping system in use, a subnet mask of 255.255.255.0 should be applied. This provides each organization with an assigned AS number of up to 255 multicast group addresses for use on the Internet.

For more information on either the practice of global scoping for multicast address allocation or the use and allocation of AS numbers, use the IETF Web site. For global scoping, refer directly to the current available draft entitled "Static Allocation in 233/8," and for more information on AS numbering, refer to RFC 1930, "Guidelines for creation, selection, and registration of an Autonomous System (AS)."

MADCAP and DHCP

The Windows Server 2003 DHCP service supports both the DHCP and MADCAP protocols. These protocols function separately and are not dependent on each other. For example, a DHCP client might or might not be a MADCAP client and a MADCAP client might or might not be a DHCP client.

It is also worth mentioning that the DHCP Server service can be used to deploy MADCAP servers independent of how DHCP servers are used on your network. For example, to install Windows Server 2003 DHCP for MADCAP service only:

1. Create multicast scopes.
2. Do not create other scopes or superscopes.

Only where other scopes or superscopes are configured does the server computer also function as a DHCP server.

For more information about MADCAP, see Multicast address allocation.

Important

- Multicast scopes and MADCAP only provide a mechanism for dynamically allocating IP address configuration for multicast-ranged IP addresses. Other network configuration details are normally required to enable multicasting for your deployment needs. For more information, see Checklist: Installing IP multicast video conferencing.

Notes

- Multicast scopes do not require or support the use of DHCP options, but can be configured with a finite lifetime, enabling the multicast scope to expire and be removed by the server.
- Web addresses can change, so you might be unable to connect to the Web site or sites mentioned here.
- For more information, see Manage Multicast Scopes.

Integrating DNS with DHCP

Updated: March 28, 2003

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Windows Server 2003 DNS supports DHCP by means of the dynamic update of DNS zones. By integrating DHCP and DNS in a DNS deployment, you can provide your network resources with dynamic addressing information stored in DNS. To enable this integration, you can use the Windows Server 2003 DHCP service.

The dynamic update standard, specified in RFC 2136: *Dynamic Updates in the Domain Name System (DNS UPDATE)*, automatically updates DNS records. Both Windows Server 2003 and Windows 2000 support dynamic update, and both clients and DHCP servers can send dynamic updates when their IP addresses change.

Dynamic update enables a DHCP server to register address (A) and pointer (PTR) resource records on behalf of a DHCP client by using DHCP Client FQDN option 81. Option 81 enables the DHCP client to provide its FQDN to the DHCP server. The DHCP client also provides instructions to the DHCP server describing how to process DNS dynamic updates on behalf of the DHCP client.

The DHCP server can dynamically update DNS A and PTR records on behalf of DHCP clients that are not capable of sending option 81 to the DHCP server. You can also configure the DHCP server to discard client A and PTR records when the DHCP client lease is deleted. This reduces the time needed to manage these records manually and provides support for DHCP clients that cannot perform dynamic updates. In addition, dynamic update simplifies the setup of Active Directory by enabling domain controllers to dynamically register SRV resource records.

If the DHCP server is configured to perform DNS dynamic updates, it performs one of the following actions:

- The DHCP server updates resource records at the request of the client. The client requests the DHCP server to update the DNS PTR record on behalf of the client, and the client registers A.
- The DHCP server updates DNS A and PTR records regardless of whether the client requests this action or not.

By itself, dynamic update is not secure because any client can modify DNS records. To secure dynamic updates, you can use the secure dynamic update feature provided in Windows Server 2003. To delete outdated records, you can use the DNS server aging and scavenging feature.

Active Directory

From Wikipedia, the free encyclopedia

Jump to: navigation, search

This article includes a list of references, but **its sources remain unclear because it has insufficient inline citations**.
Please help to improve this article by introducing more precise citations where appropriate.
(December 2008)

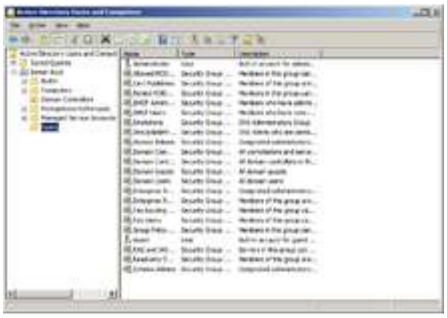

Typically Active Directory is managed using the graphical Microsoft Management Console.

Active Directory is a technology created by Microsoft that provides a variety of network services, including:

- Lightweight Directory Access Protocol LDAP is the industry standard directory access protocol, making Active Directory widely accessible to management and query applications. Active Directory supports LDAPv3 and LDAPv2.
- Kerberos-based authentication
- DNS-based naming and other network information
- Central location for network administration and delegation of authority [2]
- Information security and single sign-on for user access to networked based resources [3]
- The ability to scale up or down easily [4]
- Central storage location for application data [5]
- Synchronization of directory updates amongst several servers [6]

Using the same database, for use primarily in Windows environments, Active Directory also allows administrators to assign policies, deploy software, and apply critical updates to an organization. Active Directory stores information and settings in a central database. Active Directory networks can vary from a small installation with a few computers, users and printers to tens of thousands of users, many different domains and large server farms spanning many geographical locations.

Active Directory was previewed in 1999, released first with Windows 2000 Server edition, and revised to extend functionality and improve administration in Windows Server 2003. Additional improvements were made in Windows Server 2003 R2. Active Directory was refined further in Windows Server 2008 and Windows Server 2008 R2 and was renamed **Active Directory Domain Services**.

Active Directory was called **NTDS** (NT Directory Service) in older Microsoft documents. This name can still be seen in some Active Directory binaries.

Contents

[hide]

[edit] Structure

[edit] Objects

An Active Directory structure is a **hierarchical framework** of objects. The objects fall into two broad categories: resources (e.g., printers) and security principals (user or computer accounts and groups). Security principals are Active Directory objects that are assigned unique security identifiers (SIDs) used to control access and set security.

Each object represents a single entity — whether a user, a computer, a printer, or a group — and its attributes. Certain objects can also be containers of other objects. An object is uniquely identified by its name and has a set of attributes — the characteristics and information that the object can contain — defined by a schema, which also determines the kinds of objects that can be stored in Active Directory.

Each attribute object can be used in several different schema class objects. The schema object exists to allow the schema to be extended or modified when necessary. However, because each schema object is integral to the definition of Active Directory objects, deactivating or changing these objects can have serious consequences because it will fundamentally change the structure of Active Directory itself. A schema object, when altered, will automatically propagate through Active Directory and once it is created it can only be deactivated — not deleted. Changing the schema usually requires a fair amount of planning.[1]

[edit] Sites

A **Site** object in Active Directory represents a geographic location that hosts networks. Sites contain objects called subnets.[2] Sites can be used to assign Group Policy Objects, facilitate the discovery of resources, manage active directory replication, and manage network link traffic. Sites can be linked to other Sites. Site-linked objects may be assigned a cost value that represents the speed, reliability, availability, or other real property of a physical resource. Site Links may also be assigned a schedule.

[edit] Forests, trees, and domains

All objects inside a common directory database is known as a domain. Each domain stores information only about the objects that belong to that domain. A tree consists of a single domain or multiple domains in a contiguous namespace. A forest is a collection of trees and represents the outermost boundary within which users, computers, groups, and other objects exist. The forest is the security boundary for Active Directory.

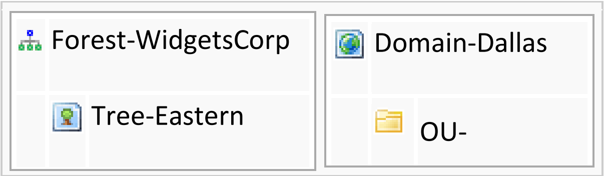

- Domain-Boston
- Domain-NewYork
- Domain-Philly
- Tree-Southern
  - Domain-Atlanta
  - Domain-Dallas

- Marketing
  - Donn
  - Mark
  - Steve
- OU-Sales
  - Bill
  - Ralph

Example of the geographical organizing of zones of interest within trees and domains.

The Active Directory framework that holds the objects can be viewed at a number of levels. At the top of the structure is the *forest.* A forest is a collection of multiple trees that share a common global catalog, directory schema, logical structure, and directory configuration. The forest, tree, and domain are the logical parts in an Active Directory network.

The Active Directory forest contains one or more transitive, trust-linked *trees*. A tree is a collection of one or more *domains* and domain trees in a contiguous namespace, again linked in a transitive trust hierarchy. Domains are identified by their DNS name structure, the namespace.

[edit] Flat-filed, simulated hierarchy

The objects held within a domain can be grouped into containers called Organizational Units (OUs).[3] OUs give a domain a hierarchy, ease its administration, and can give a resemblance of the structure of the organization in organizational or geographical terms. OUs can contain OUs – indeed, domains are containers in this sense – and can hold multiple nested OUs. Microsoft recommends as few domains as possible in Active Directory and a reliance on OUs to produce structure and improve the implementation of policies and administration. The OU is the common level at which to apply group policies, which are Active Directory objects themselves called Group Policy Objects (GPOs), although policies can also be applied to domains or sites (see below). The OU is the level at which administrative powers are commonly delegated, but granular delegation can be performed on individual objects or attributes as well.

However, Organizational Units are just an abstraction for the administrator, and do not function as true containers; the underlying domain operates as if objects were all created in a simple flat-file structure, without any OUs. It is not possible, for example, to create two user accounts with an identical username (sAMAccountName) in two separate OUs, such as "fred.staff-ou.domain" and "fred.student-ou.domain". This is so because sAMAccountName, a user object attribute, is constrained to be unique across the entire domain. However, note that you can have two different "Freds" within AD, each with his own different account name (sAMAccountName) - for e.g. Fred Smith (fsmith), and Fred A. Smith (fasmith). Note that they are both Fred Smiths, albeit one with a middle initial, they have different account names (sAMAccountName) - fsmith, and fasmith.

By contrast, there are other vendor directories such as Novell eDirectory that allow naming attribute duplication across separate OUs. Each user logs in by specifying the *context* of their account, which is similar to the *current working directory* of a file system. Context normally operates in relative form: if the login prompt context is "staff-ou.accounts-ou.organization", people with accounts in that specific OU need only type their username "fred". But if the login prompt context were set to be one level higher, at "accounts-ou.organization", people would need to specify the OU within that context: "fred.staff-ou". Context can also be specified in *absolute* form similar to an *absolute directory path* by using a leading period: ".fred.staff-ou.accounts-ou.organization", which disregards the current login prompt context.

Novell additionally provides login prompt functionality known as *contextless login*[4] to permit searching the directory structure via LDAP for all possible matching or similar usernames, making the Novell login process operate similar to Microsoft's flat-file structure that searches the entire domain for accounts regardless of the account's location in the OUs. The concept of account context in the directory does not apply to Active Directory, since object name duplication within a single domain is not permitted to occur in the first place.

Because duplicate usernames cannot exist within separate OUs of a single active directory domain, unique account name generation poses a significant challenge for organizations with hundreds to thousands of users that are part of a generalized mass that can not be easily subdivided into separate domains, such as students in a public school system or university that must be able to login on any computer across the district buildings or campus network.

As the number of users in a domain increases, simple username creation methods such as "first initial, middle initial, last name" will fail due to having so many common names like *Smith* or *Johnson* in the collective mass that result in having duplications, such as two *JASmith*, which requires randomly adding a number to the end *(JASmith1)* to further differentiate it for one of the two people. At some point of increasingly many users and name duplications, the network IT staff may give up on attempts at making usernames personally memorable, and the username simply becomes a serial number 5 to 10 digits long to provide sufficient naming uniqueness within a single domain.

[edit] Shadow Groups

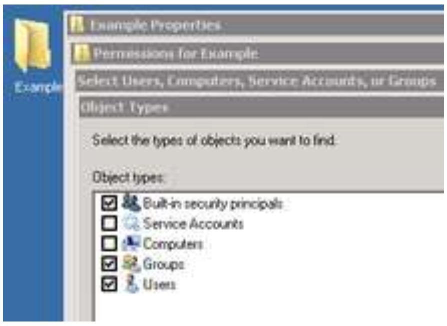

Select the types of objects you want to find.

Object types:

☑ Built-in security principals
☐ Service Accounts
☐ Computers
☑ Groups
☑ Users

In Active Directory, organizational units can not be assigned as owners or trustees. Only groups are selectable, and members of OUs can not be collectively assigned rights to directory objects.

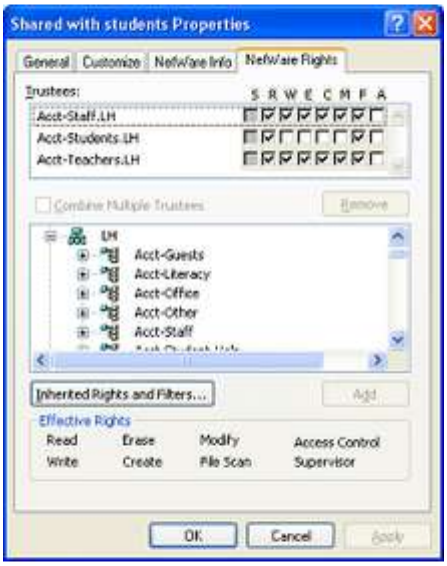

Trustees:

	S	R	W	E	C	M	F	A
Acct-Staff.LH	□	☑	☑	☑	☑	☑	☑	□
Acct-Students.LH	□	☑	□	□	□	□	☑	□
Acct-Teachers.LH	□	☑	☑	☑	☑	☑	☑	□

Effective Rights
Read Erase Modify Access Control
Write Create File Scan Supervisor

Unlike Active Directory, Novell eDirectory allows organizational units and all users within the OU to be assigned rights to an object, without having to create shadow groups representing the users in each OU.

It is often useful to associate a collection of users to all share access rights to particular file or secured resource, but with Active Directory it is not possible to choose an OU containing all users that need rights. A user group can be selected to accomplish this, but all users within a particular OU are not automatically made members of a group representing that OU.

Groups can be manually created to duplicate the account membership structure within OUs, but it is an extra step of the account creation process by the administrator to remember all the various groups each new user needs to join. If the administrator forgets this manual step, the users will experience problems until the group memberships are corrected.

To make up for this non-automated deficiency, network administrators can write their own custom scripts which periodically run on the server and use LDAP access commands to add or remove users from groups representing the OUs of the users, known as *Shadow Groups*. Microsoft refers to shadow groups in the Server 2008 Reference documentation, but does not explain how to create them.[5] Once created, these shadow groups are selectable in place of the OU in the administrative console tools.

The naming of shadow groups is complicated by the fact that OUs can be nested but groups cannot. Groups can only exist in the root of the domain, and group names are limited in length so matching the naming of a deeply nested string of OUs for a very large domain is difficult.

Novell eDirectory supports the creation of user groups, but OUs can be natively selected as the assigned owner of a secured resource, so shadow groups are unnecessary.

[edit] Structural divisions to improve performance

Active Directory also supports the creation of *Sites*, which are physical, rather than logical, groupings defined by one or more IP subnets.[6] Sites distinguish between locations connected by low-speed (e.g., WAN, VPN) and high-speed (e.g., LAN) connections. Sites are independent of the domain and OU structure and are common across the entire forest. Sites are used to control network traffic generated by replication and also to refer clients to the nearest domain controllers. Exchange 2007 also uses the site topology for mail routing. Policies can also be applied at the site level.

The actual division of an organization's information infrastructure into a hierarchy of one or more domains and top-level OUs is a key decision. Common models are by business unit, by geographical location, by IT Service, or by object type. These models are also often used in combination. OUs should be structured primarily to facilitate administrative delegation, and secondarily, to facilitate group policy application. Although OUs form an administrative boundary, the only true security boundary is the forest itself[7] and an administrator of any domain in the forest must be trusted across all domains in the forest.

Physically the Active Directory information is held on one or more equal peer domain controllers (DCs), replacing the NT PDC/BDC model. Each DC has a copy of the Active Directory; changes on one computer being synchronized (converged) between all the DC computers by *multi-master replication*. Servers joined to Active Directory that are not domain controllers are called Member Servers.[8]

The Active Directory database is split into different stores or *partitions*.[9] Microsoft often refers to these partitions as 'naming contexts'.[10] The 'Schema' partition contains the definition of object classes and attributes within the Forest. The 'Configuration' partition contains information on the physical structure and configuration of the forest (such as the site topology). The 'Domain' partition holds all objects created in that domain. The first two partitions replicate to all domain controllers in the Forest. The Domain partition replicates only to Domain Controllers within its domain. A subset of objects in the domain partition is also replicated to domain controllers that are configured as global catalogs.

Unlike earlier versions of Windows, which used NetBIOS to communicate, Active Directory is fully integrated with DNS and TCP/IP—DNS is *required*. To be fully functional, the DNS server must support SRV resource records or service records.

Active Directory replication is 'pull' rather than 'push'.[11] The *Knowledge Consistency Checker* (KCC) creates a replication topology of *site links* using the defined *sites* to manage traffic. Intrasite replication is frequent and automatic as a result of change notification, which triggers peers to begin a pull replication cycle. Intersite replication intervals are less frequent and do not use change notification by default, although this is configurable and can be made identical to intrasite replication. A different 'cost' can be given to each link (e.g., DS3, T1, ISDN etc.) and the site link topology will be altered accordingly by the KCC. Replication between domain controllers may occur transitively through several site links on same-protocol *site link bridges*, if the cost is low, although KCC automatically costs a direct site-to-site link lower than transitive connections. Site-to-site replication can be configured to occur between a *bridgehead server* in each site, which then replicates the changes to other DCs within the site.

In a multi-domain forest the Active Directory database becomes partitioned. That is, each domain maintains a list of only those objects that belong in that domain. So, for example, a user created in Domain A would be listed only in Domain A's domain controllers. Global catalog (GC) servers are used to provide a global listing of all objects in the Forest.[12] The Global catalog is held on domain controllers configured as global catalog servers. Global Catalog servers replicate to themselves all objects from all domains and hence, provide a global listing of objects in the forest. However, in order to minimize replication traffic and to keep the GC's database small, only selected attributes of each object are replicated. This is called the partial attribute set (PAS). The PAS can be modified by modifying the schema and marking attributes for replication to the GC.

Replication of Active Directory uses Remote Procedure Calls (RPC over IP [RPC/IP]). Between Sites you can also choose to use SMTP for replication, but only for changes in the Schema, Configuration, or Partial Attribute Set (Global Catalog) NCs. SMTP cannot be used for replicating the default Domain partition.[13]

The Active Directory database, the *directory store*, in Windows 2000 Server uses the JET Blue-based Extensible Storage Engine (ESE98), limited to 16 terabytes and 1 billion objects in each domain controller's database. Microsoft has created NTDS databases with more than 2 billion objects.[*citation needed*] (NT4's Security Account Manager could support no more than 40,000 objects). Called NTDS.DIT, it has two main tables: the *data table* and the *link table*. In Windows Server 2003 a third main table was added for security descriptor single instancing.[14]

The features of Active Directory may be accessed programmatically via the COM interfaces provided by **Active Directory Service Interfaces**.[15]

Active Directory is a necessary component for many Windows services in an organization such as Exchange, Security.

[edit] FSMO Roles

Flexible Single Master Operations (**FSMO**, sometimes pronounced "fizz-mo") roles are also known as operations master roles. Although the AD domain controllers operate in a multi-master model, i.e. updates can occur in multiple places at once, there are several roles that are necessarily single instance:

Role Name	Scope	Description
Schema Master	1 per forest	Controls and handles updates/modifications to the Active Directory schema.
Domain Naming	1 per forest	Controls the addition and removal of domains from the forest if

Master		present in root domain
PDC Emulator	1 per domain	Provides backwards compatibility for NT4 clients for PDC operations (like password changes). The PDCs also run domain specific processes such as the Security Descriptor Propagator (SDPROP), and is the master time server within the domain. It also handles external trusts, the DFS consistency check, holds the most current passwords and manages all GPOs as default server.
RID Master	1 per domain	Allocates pools of unique identifier to domain controllers for use when creating objects
Infrastructure Master	1 per domain/partition	Synchronizes cross-domain group membership changes. The infrastructure master cannot run on a global catalog server (GCS)(unless all DCs are also GCs, or environment consists of a single domain)

[edit] Trust

To allow users in one domain to access resources in another, Active Directory uses trusts.[16] Trusts inside a forest are automatically created when domains are created. The forest sets the default boundaries of trust, not the domain, and implicit, transitive trust is automatic for all domains within a forest. As well as two-way transitive trust, AD trusts can be a *shortcut* (joins two domains in different trees, transitive, one- or two-way), *forest* (transitive, one- or two-way), *realm* (transitive or nontransitive, one- or two-way), or *external* (nontransitive, one- or two-way) in order to connect to other forests or non-AD domains.[17]

[edit] Trusts in Windows 2000 (native mode)

- **One-way trust** – One domain allows access to users on another domain, but the other domain does not allow access to users on the first domain.
- **Two-way trust** – Two domains allows access to users on both domains.
- **Trusting domain** – The domain that allows access to users from a trusted domain.
- **Trusted domain** – The domain that is trusted; whose users have access to the trusting domain.
- **Transitive trust** – A trust that can extend beyond two domains to other trusted domains in the forest.
- **Intransitive trust** – A one way trust that does not extend beyond two domains.
- **Explicit trust** – A trust that an admin creates. It is not transitive and is one way only.
- **Cross-link trust** – An explicit trust between domains in different trees or in the same tree when a descendant/ancestor (child/parent) relationship does not exist between the two domains.

Windows 2000 Server – supports the following types of trusts:

- Two-way transitive trusts.

- One-way intransitive trusts.

Additional trusts can be created by administrators. These trusts can be:

- Shortcut

Windows Server 2003 offers a new trust type – the forest root trust. This type of trust can be used to connect Windows Server 2003 forests if they are operating at the 2003 forest functional level. Authentication across this type of trust is Kerberos based (as opposed to NTLM). Forest trusts are also transitive for all the domains in the forests that are trusted. Forest trusts, however, are not transitive.

[edit] ADAM/AD LDS

Active Directory Application Mode (ADAM) is a light-weight implementation of Active Directory. ADAM is capable of running as a service, on computers running Microsoft Windows Server 2003 or Windows XP Professional. ADAM shares the code base with Active Directory and provides the same functionality as Active Directory, including an identical API, but does not require the creation of domains or domain controllers.

Like Active Directory, ADAM provides a *Data Store*, which is a hierarchical datastore for storage of directory data, a *Directory Service* with an LDAP *Directory Service Interface*. Unlike Active Directory, however, multiple ADAM instances can be run on the same server, with each instance having its own and required by applications making use of the ADAM directory service.

In Windows Server 2008, ADAM has been renamed AD LDS (Lightweight Directory Services).[18]

[edit] Integrating Unix into Active Directory

Varying levels of interoperability with Active Directory can be achieved on most Unix-like operating systems through standards compliant LDAP clients, but these systems usually lack the automatic interpretation of many attributes associated with Windows components, such as Group Policy and support for one-way trusts.

There are also third-party vendors who offer Active Directory integration for Unix platforms (including UNIX, Linux, Mac OS X, and a number of Java- and UNIX-based applications). Some of these vendors include Centrify (DirectControl), Computer Associates (UNAB), CyberSafe Limited (TrustBroker), Likewise Software (Open or Enterprise), Quest Software (Authentication Services) and Thursby Software Systems (ADmitMac). The open source Samba software provides a way to interface with Active Directory and join the AD domain to provide authentication and authorization: version 4 (in alpha as of October 2009) can act as a peer Active Directory domain controller.[19] Microsoft is also in this market with their free Microsoft Windows Services for UNIX product.

The schema additions shipped with Windows Server 2003 R2 include attributes that map closely enough to RFC 2307 to be generally usable. The reference implementation of RFC

2307, nss_ldap and pam_ldap provided by PADL.com, contains support for using these attributes directly, provided they have been populated. The default Active Directory schema for group membership complies with the proposed extension, RFC 2307bis. Windows Server 2003 R2 includes a Microsoft Management Console snap-in that creates and edits the attributes.

An alternate option is to use another directory service such as 389 Directory Server (formerly Fedora Directory Server) or Sun Microsystems Sun Java System Directory Server, which can perform a two-way synchronization with Active Directory and thus provide a "deflected" integration with Active Directory as Unix and Linux clients will authenticate to FDS and Windows Clients will authenticate to Active Directory. Another option is to use OpenLDAP with its translucent overlay, which can extend entries in any remote LDAP server with additional attributes stored in a local database. Clients pointed at the local database will see entries containing both the remote and local attributes, while the remote database remains completely untouched.

[edit] See also

- FreeIPA
- Active Directory Explorer
- Directory Services Restore Mode
- Flexible single master operation
- List of LDAP software
- AGDLP (implementing role based access controls using nested groups)

[edit] Notes

1. ^ *Windows Server 2003: Active Directory Infrastructure*. Microsoft Press. 2003. pp. 1–8 – 1–9. ISBN 0-7356-1438-5.
2. ^ "Managing Sites". Microsoft Corporation-0. http://technet.microsoft.com/en-us/library/bb727051.aspx. "An Active Directory site object represents a collection of Internet Protocol (IP) subnets, usually constituting a physical Local Area Network (LAN)."
3. ^ "Organizational Units". Microsoft Corporation. 2010. http://technet.microsoft.com/en-us/library/cc978003.aspx. "An organizational unit in Active Directory is analogous to a directory in the file system"
4. ^ Novell: Taking Things Out of Context: Using LDAP Contextless Login in Your Network, 01 September 2003 [1]
5. ^ Microsoft Server 2008 Reference refers to "shadow groups" but does not explain how to create them. http://technet.microsoft.com/en-us/library/cc770394%28WS.10%29.aspx
6. ^ "Sites overview". Microsoft Corporation. 2005-01-21. http://technet.microsoft.com/en-us/library/cc782048(WS.10).aspx. "A site is a set of well-connected subnets."
7. ^ "Specifying Security and Administrative Boundaries". Microsoft Corporation. 2005-01-23. http://technet.microsoft.com/en-us/library/cc755979(WS.10).aspx. "However, service administrators have abilities that cross domain boundaries. For this reason, the forest is the ultimate security boundary, not the domain."
8. ^ "Planning for domain controllers and member servers". Microsoft Corporation. 2005-01-21. http://technet.microsoft.com/en-us/library/cc737059(WS.10).aspx. "[...] member servers, [...] belong to a domain but do not contain a copy of the Active Directory data."

9. ^ "Directory data store". Microsoft Corporation. 2005-01-21. http://technet.microsoft.com/en-us/library/cc736627(WS.10).aspx. "Active Directory uses four distinct directory partition types to store [...] data. Directory partitions contain domain, configuration, schema, and application data."
10. ^ Andreas Luther. "Active Directory Replication Traffic". Microsoft Corporation. http://technet.microsoft.com/en-us/library/bb742457.aspx. Retrieved 2010-05-26. "The Active Directory is made up of one or more naming contexts or partitions."
11. ^ "What Is the Active Directory Replication Model?". Microsoft Corporation. 2003-03-28. http://technet.microsoft.com/en-us/library/cc737314(WS.10).aspx. "Domain controllers request (pull) changes rather than send (push) changes that might not be needed."
12. ^ "What Is the Global Catalog?". Microsoft Corporation. 2009-12-10. http://technet.microsoft.com/en-us/library/cc728188(WS.10).aspx. "[...] a domain controller can locate only the objects in its domain. [...] The global catalog provides the ability to locate objects from any domain [...]"
13. ^ "What Is Active Directory Replication Topology?". Microsoft Corporation. 2003-03-28. http://technet.microsoft.com/en-us/library/cc775549(WS.10).aspx. "SMTP can be used to transport nondomain replication [...]"
14. ^ Large AD database? Probably not this large...
15. ^ Active Directory Service Interfaces, Microsoft
16. ^ "Domain and Forest Trusts Technical

Managing DHCP, Windows Internet Name Service, and Internet Authentication Service

A network administrator might use SNMP to assist in the following duties:

- Viewing and changing parameters in the LAN Manager and MIB-II MIBs.
- Monitoring and configuring parameters for any WINS servers on the network.
- Monitoring DHCP servers.
- Using System Monitor to monitor TCP/IP- related performance counters (Internet Control Message Protocol (ICMP), IP, Network Interface, TCP, UDP, DHCP, FTP, WINS, and IIS performance counters).

For more information about System Monitor, see the Microsoft® Windows® 2000 Professional Resource Kit.

Use the tools on the *Windows 2000 Resource Kit* companion CD to perform simple SNMP management functions.

Using System Monitor

All System Monitor counters installed on a computer can be viewed with SNMP. To view System Monitor counters with SNMP, use the Perf2MIB tool provided on the *Windows 2000 Resource Kit* companion CD. For additional information about how to use the Perf2mib.exe tool, see Tools Help on the companion CD.

Top Of Page

Managing DHCP

The Windows 2000 – based DHCP server objects and IIS objects can be monitored but not configured by using SNMP .

Top Of Page

Managing WINS

All but a few of the WINS server objects can be monitored and configured by using SNMP. For information about what WINS parameters can be configured using SNMP, see "MIB Object Types" in this book. Any WINS objects defined with read/write permissions can be configured.

Top Of Page

Managing IAS

Internet Authentication Server (IAS) implements the RADIUS authentication and accounting MIBs, which permit IAS objects to be monitored and configured using SNMP. Any IAS objects defined with read/write permissions can be configured

Network Infrastructure encompassing

Remote Access Troubleshooting

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Troubleshooting

What problem are you having?

- Unable to establish a remote access connection.
- Unable to access resources beyond the remote access server.
- Callback is not working.
- Not sure how to configure the remote access server to make as well as receive remote access connections (dial out and dial in).

Unable to establish a remote access connection.
Cause: Your modem is not working properly.

Solution: Verify that your modem is working properly.

See also: Troubleshooting modems

Cause: The Routing and Remote Access service is not started on the remote access server.

Solution: Verify the state of the Routing and Remote Access service on the remote access server.

See also: Monitor the Routing and Remote Access service

Cause: Remote access is not enabled on the remote access server.

Solution: Enable remote access on the remote access server.

See also: Enable the remote access server

Cause: Dial-in, PPTP, or L2TP ports are not enabled for inbound remote access connections.

Solution: Enable PPTP, L2TP, or dial-in ports for inbound remote access connections, as needed.

See also: Configure ports for remote access

Cause: The LAN protocols being used by the remote access clients are not configured on the remote access server to allow remote access connections.

Solution: Verify that remote access connections are allowed on the **IP** or **AppleTalk** tabs on the properties of a server. If you do not have one or more of these tabs, then the corresponding network protocol is not installed on the server.

See also: View properties of the remote access server

Cause: All of the remote access ports on the remote access server are already being used by currently connected remote access clients or demand-dial routers.

Solution: You can verify whether all of the remote access ports on the remote access server are not already being used by clicking **Ports** in Routing and Remote Access. If all ports are busy, disconnect any inactive connections, or consider adding more ports.

Cause: The remote access client and the remote access server in conjunction with a remote access policy are not configured to use at least one common authentication method.

Solution: Configure the remote access client and the remote access server in conjunction with a remote access policy to use at least one common authentication method.

If a remote access client running Windows 95 is attempting a dial-up connection, verify that the remote access server is not requiring only MS-CHAP v2 authentication. A remote access client running Windows 95 supports the use of MS-CHAP v2 over virtual private network (VPN) connections, not the use of MS-CHAP v2 over dial-up connections.

Cause: The remote access client and the remote access server in conjunction with a remote access policy are not configured to use at least one common encryption strength.

Solution: Configure the remote access client and the remote access server in conjunction with a remote access policy to use at least one common encryption strength.

See also: Configure encryption

Cause: The remote access connection does not have the appropriate permissions through dial-in properties of the user account and remote access policies.

Solution: Verify that the remote access connection has the appropriate permissions through dial-in properties of the user account and remote access policies.

In order for the connection to be established, the settings of the connection attempt must:

- Match all of the conditions of at least one remote access policy.
- Be granted remote access permission through the local user account (set to **Allow access**), or through the domain user account (set to **Control access through Remote Access Policy**) and the remote access permission of the matching remote access policy (set to **Grant remote access permission**).
- Match all the settings of the profile.
- Match all the settings of the dial-in properties of the user account.

See also: Introduction to remote access policies; Accepting a connection attempt

Cause: The settings of the remote access policy profile are in conflict with properties of the remote access server.

Solution: Verify that the settings of the remote access policy profile are not in conflict with properties of the remote access server.

The properties of the remote access policy profile and the properties of the remote access server both contain settings for:

- Multilink
- Bandwidth allocation protocol
- Authentication protocols

If the settings of the profile of the matching remote access policy are in conflict with the settings of the remote access server, the connection attempt is rejected. For example, if the matching remote access policy profile specifies that the EAP-TLS authentication protocol must be used and EAP is not enabled on the remote access server, the connection attempt is rejected.

See also: Enable authentication protocols; Configure authentication

Cause: The credentials of the remote access client (user name, password, and domain name) are incorrect and cannot be validated by the remote access server.

Solution: Verify that the credentials of the remote access client (user name, password, and domain name) are correct and can be validated by the remote access server.

Cause: There are not enough addresses in the static IP address pool.

Solution: If the remote access server is configured with a static IP address pool, verify that there are enough addresses in the pool. If all of the addresses in the static pool have been allocated to connected remote access clients, the remote access server is unable to allocate an IP address, and the connection attempt is rejected. Modify the static IP address pool if needed.

See also: Create a static IP address pool

Cause: The authentication provider of the remote access server is improperly configured.

Solution: Verify the configuration of the authentication provider. You can configure the remote access server to use either Windows Authentication or Remote Authentication Dial-In User Service (RADIUS) to authenticate the credentials of the remote access client.

See also: Use RADIUS authentication

Cause: The remote access server cannot access Active Directory.

Solution: For a remote access server that is a member server of a domain that is configured for Windows Authentication, verify that:

- The **RAS and IAS Servers** security group exists. If not, then create the group and set the group type to **Security** and the group scope to **Domain local**.
- The **RAS and IAS Servers** security group has Read permission to the **RAS and IAS Servers Access Check** object.
- The computer account of the remote access server computer is a member of the **RAS and IAS Servers** security group. You can use the **netsh ras show registeredserver** command to view the

current registration. You can use the **netsh ras add registeredserver** command to register the server in a specified domain.

If you add or remove the remote access server to the **RAS and IAS Servers** security group, the change does not take effect immediately (due to the way that Active Directory information is cached). For the change to take effect immediately, you need to restart the remote access server computer.

- The remote access server has joined the domain.

See also: Create a new group; Verify permissions for the RAS and IAS security group; Netsh commands for remote access

Cause: A remote access server running Windows NT 4.0 cannot validate connection requests.

Solution: If remote access clients are dialing in to a remote access server running Windows NT 4.0 that is a member of a Windows 2000 mixed domain, verify that the Everyone group is added to the Pre-Windows 2000 Compatible Access group with the **net localgroup "Pre-Windows 2000 Compatible Access"** command. If not, type **net localgroup "Pre-Windows 2000 Compatible Access" everyone /add** at the command prompt on a domain controller computer and then restart the domain controller computer.

See also: Member server in a domain

Cause: The Windows Fax service is enabled and your modem does not support adaptive answer.

Solution: If the Windows Fax service and the Routing and Remote Access service are sharing the same modem, verify that the modem supports adaptive answer. If the modem does not support adaptive answer, you must disable fax on the modem to receive incoming remote access connections.

Cause: You are using MS-CHAP v1 and a user password over 14 characters long.

Solution: Either use a different authentication protocol such as MS-CHAP v2 or change your password so that it is 14 characters or less.

See also: MS-CHAP; Enable authentication protocols

Unable to access resources beyond the remote access server.
Cause: For IP-based remote access clients, IP routing is not enabled.

Solution: Verify that IP routing is enabled on the **IP** tab on the properties of the server.

See also: View properties of the remote access server

Cause: For AppleTalk-based remote access clients, network access is not allowed.

Solution: Verify that network access is allowed on the **AppleTalk** tab on the properties of a server. If you do not have the AppleTalk tab, then AppleTalk is not installed on the server.

See also: View properties of the remote access server

Cause: A static IP address pool is configured but there are no routes back to the remote access clients.

Solution: If the remote access server is configured to use a static IP address pool, verify that the routes to the ranges of addresses of the static IP address pool are reachable by the hosts and routers of the intranet. If not, then IP routes consisting of the address ranges of the static IP address pool, as defined by the IP address and mask of each range, must be added to the routers of the intranet or the routing protocol of your routed infrastructure on the remote access server must be enabled. If the routes to the remote access client subnets are not present, remote access clients cannot receive traffic from locations on the intranet. A route for the network is implemented either through static routing entries or through a routing protocol, such as Open

Shortest Path First (OSPF) or Routing Information Protocol (RIP). OSPF is not available on Windows XP 64-bit Edition (Itanium) and the 64-bit versions of the Windows Server 2003 family.

If the remote access server is configured to use DHCP for IP address allocation and no DHCP server is available, the remote access server allocates addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the remote access server is attached is also using APIPA addresses.

If the remote access server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IP addresses. By default, the remote access server randomly chooses the adapter to use to obtain IP addresses through DHCP. If there is more than one LAN adapter, then the Routing and Remote Access service may choose a LAN adapter for which there is no DHCP server available.

If the static IP address pool consists of ranges of IP addresses that are a subset of the range of IP addresses for the network to which the remote access server is attached, verify that the ranges of IP addresses of the static IP address pool are not assigned to other TCP/IP nodes, either through static configuration or through DHCP.

See also: Create a static IP address pool

Cause: Packet filters on the remote access policy profile are preventing the flow of IP traffic.

Solution: Verify that there are no configured TCP/IP packet filters on the profile properties of the remote access policies on the remote access server (or the RADIUS server if Internet Authentication Service is used) that are preventing the sending or receiving of TCP/IP traffic.

You can use remote access policies to configure TCP/IP input and output packet filters that control the exact nature of TCP/IP traffic allowed on the remote access connection. Verify that the profile TCP/IP packet filters are not preventing the flow of needed traffic.

See also: Configure IP options

Callback is not working.
Cause: Callback is improperly configured on the dial-in properties of the user account.

Solution: Verify the callback configuration on the dial-in properties of the user account.

See also: Configure caller ID and callback

Cause: **Link control protocol (LCP) extensions** is disabled on the **PPP** tab for the properties of the remote access server.

Solution: Enable **Link control protocol (LCP) extensions** on the **PPP** tab for the properties of the remote access server.

See also: View properties of the remote access server

Cause: Callback numbers may be truncated when a remote access server running Windows NT 4.0 requests dial-in properties of a user account in a Windows 2000 native or Windows Server 2003 domain.

Solution: Reconfigure callback numbers on the dial-in properties of the user account.

See also: Configure caller ID and callback

For more information on troubleshooting remote access VPN connections, see Troubleshooting remote access VPNs.

To reset Routing and Remote Access back to default values, see Reset the Routing and Remote Access service.

Not sure how to configure the remote access server to make as well as receive remote access connections (dial out and dial in).

Cause: The remote access server is either not configured as a remote access server (dial in) or not configured for demand-dial routing (dial out).

Solution: In the console tree, right-click the remote access server, and click **Properties**. On the **General** tab, make sure that both the **Router** and the **Remote access server** check boxes are selected, and make sure that the **LAN and demand-dial routing** option is selected under **Router**. If these options are not configured, configure them, and click **OK**. In the console tree, click **Network Interfaces** and make sure that you have configured demand-dial interfaces for dialing out to the servers or services to which you want to connect. If no demand-dial interfaces exist, right-click **Network Interfaces**, and click **New Demand-dial Interface**.

See also: Demand-dial routing; Routing over VPN connections; Overview of remote access

Remote Access Policy

1.0 Overview

This remote access policy defines standards for connecting to the organizational network and security standards for computers that are allowed to connect to the organizational network.

This remote access policy specifies how remote users can connect to the main organizational network and the requirements for each of their systems before they are allowed to connect. This will specify:

1. The anti-virus program remote users must use and how often it must be updated.
2. What personal firewalls they are required to run.
3. Other protection against spyware or other malware.

The remote access policy defines the methods users can use to connect remotely such as dial up or VPN. It will specify how the dial up will work such as whether the system will call the remote user back, and the authentication method. If using VPN, the VPN protocols used will be defined. Methods to deal with attacks should be considered in the design of the VPN system.

2.0 Purpose

This remote access policy is designed to prevent damage to the organizational network or computer systems and to prevent compromise or loss of data.

3.0 Approval

Any remote access using either dial-in, VPN, or any other remote access to the organizational network must be reviewed and approved by the appropriate supervisor. All employees by default will have account settings set to deny remote access. Only upon approval will the account settings be changed to allow remote access.

4.0 Remote Computer Requirements

1. The anti-virus product called ______________ is required to be operating on the computer at all times in real time protection mode.
    1. The anti-virus product shall be operated in real time on the computer. The product shall be configured for real time protection.
    2. The anti-virus library definitions shall be updated at least once per day.
    3. Anti-virus scans shall be done a minimum of once per week.

    No one should be able to stop anti-virus definition updates and anti-virus scans except for domain administrators.

2. The computer must be protected by a firewall at all times when it is connected to the internet. Acceptable products include ________________. Several popular choices include Zone Alarm, the Windows XP firewall, and Norton Personal firewall.

5.0 Remote Connection Requirements

The remote user shall use either dial-In or virtual private networking (VPN). Dial-In is typically used when the user in in a local calling area. VPN is typically used when the user would need to dial a long distance number to connect with a dial-in connection. VPN uses a local connection to an internet service provider (ISP) and creates a tunnel through the local ISP connection to the organizational network. This section specifies the requirements for Dial-In and VPN connections.

5.1 Dial-In Requirements

1. Number check - The dial in settings shall be set to perform one or the other of:
    1. Verify Caller ID to a specific number - Use this option if caller ID is available
    2. Always Call back to a specific number - If the user must connect from a location other than their designated location such as their home, they should use VPN.
2. Client Check - A requirement that must be set for Dial-In clients is that a firewall must be installed and operational. If the Dial-In client does not meet the criteria, either the connection is not allowed or the client can only access a limited area where they can get the software needed to meet the requirement.
3. Authentication - For authentication of the user, the dial in connection shall use one of:
    1. MS-CHAP version 2
    2. EAP-RADIUS

    3. EAP-TLS
    4. EAP-MD5-Challenge
4. Connection Encryption - This requirement will depend on the data you expect the remote user to be transmitting over the dial-in connection. Typically this should be encrypted especially if the user works for the Finance or Personnel department. The connection shall use one of the following encryption mechanisms:
    1. Microsoft Point to Point Encryption (MPPE)
    2. IPSec

5.2 VPN Requirements

1. Client Check - A requirement that must be set for VPN clients is that a firewall must be installed and operational. Also Anti-virus software must be installed and operational. If the VPN client does not meet the criteria, either the connection is not allowed or the client can only access a limited area where they can get the software needed to meet the requirement.
2. The connection choices are PPTP, L2TP, IPSec, and SSL. The connection shall use IPSec which encrypts the data sent through the connection.
3. Authentication - For authentication of the user, the dial in connection shall use Internet Key Exchange (IKE) with digital certificates. The other choice is Internet Key Exchange (IKE) with a preshared key.

Remote access policy

From Wikipedia, the free encyclopedia
Jump to: navigation, search

Remote access policy is a document which outlines and defines acceptable methods of remotely connecting to the internal network. It is essential in large organization where networks are geographically dispersed and extend into insecure network locations such as public networks or unmanaged home networks. It should cover all available methods to remotely access internal resources:

- dial-in (SLIP, PPP)
- ISDN/Frame Relay
- telnet access from Internet
- Cable modem

This remote access policy defines standards for connecting to the organizational network and security standards for computers that are allowed to connect to the organizational network.

This remote access policy specifies how remote users can connect to the main organizational network and the requirements for each of their systems before they are allowed to connect. This will specify:

Remote Access VPN Business Scenarios

Table Of Contents

Remote Access VPN Business Scenarios

This chapter explains the basic tasks for configuring an IP-based, remote access Virtual Private Network (VPN) on a Cisco 7200 series router. In the remote access VPN business scenario, a remote user running VPN client software on a PC establishes a connection to the headquarters Cisco 7200 series router.

The configurations in this chapter utilize a Cisco 7200 series router. If you have a Cisco 2600 series router or a Cisco 3600 series router, your configurations will differ slightly, most notably in the port slot numbering. Please refer to your model configuration guide for detailed configuration information. Please refer to the "Obtaining Documentation" section on page xii for instructions about locating product documentation.

Note In this Guide, the term `Cisco 7200 series router' implies that an Integrated Service Adaptor (ISA) or a VAM (VAM, VAM2, or VAM2+) is installed in the Cisco 7200 series router.

This chapter describes basic features and configurations used in a remote access VPN scenario. Some Cisco IOS security software features not described in this document can be used to increase performance and scalability of your VPN. For up-to-date Cisco IOS security software features documentation, refer to the *Cisco IOS Security Configuration Guide* and the *Cisco IOS Security Command Reference* for your Cisco IOS Release. To access these documents, see "Related Documentation" section on page xi.

This chapter includes the following sections:

- Scenario Description
- Configuring a Cisco IOS VPN Gateway for Use with Cisco Secure VPN Client Software
- Configuring a Cisco IOS VPN Gateway for Use with Microsoft Dial-Up Networking
- Configuring Cisco IOS Firewall Authentication Proxy
- Comprehensive Configuration Examples

Note Throughout this chapter, there are numerous configuration examples and sample configuration outputs that include unusable IP addresses. Be sure to use your own IP addresses when configuring your Cisco 7200 series router.

Scenario Description

Figure 4-1 shows a headquarters network providing a remote user access to the corporate intranet. In this scenario, the headquarters and remote user are connected through a secure tunnel that is established over an IP infrastructure (the Internet). The remote user is able to access internal, private web pages and perform various IP-based network tasks.

Figure 4-1 Remote Access VPN Business Scenario

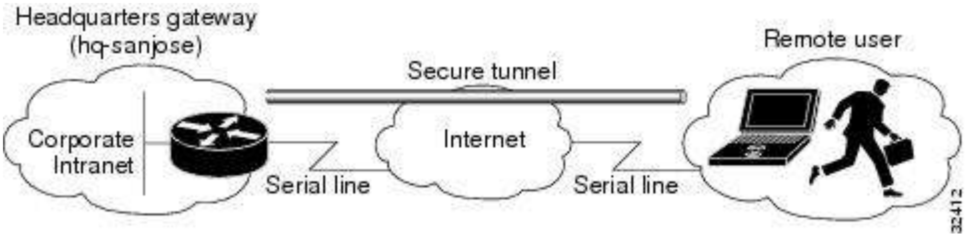

[Figure 4-2](#) shows the physical elements of the scenario. The Internet provides the core interconnecting fabric between the headquarters and remote user. The headquarters is using a Cisco IOS VPN gateway (Cisco 7200 series with an Integrated Service Adaptor (ISA) or VAM, a Cisco 2600 seriesrouter or a 3600 series router), and the remote user is running VPN client software on a PC.

The tunnel is configured on the first serial interface in chassis slot 1 (serial 1/0) of the headquarters and remote office routers. Fast Ethernet interface 0/0 of the headquarters router is connected to a corporate server and Fast Ethernet interface 0/1 is connected to a web server.

Figure 4-2 Remote Access VPN Scenario Physical Elements

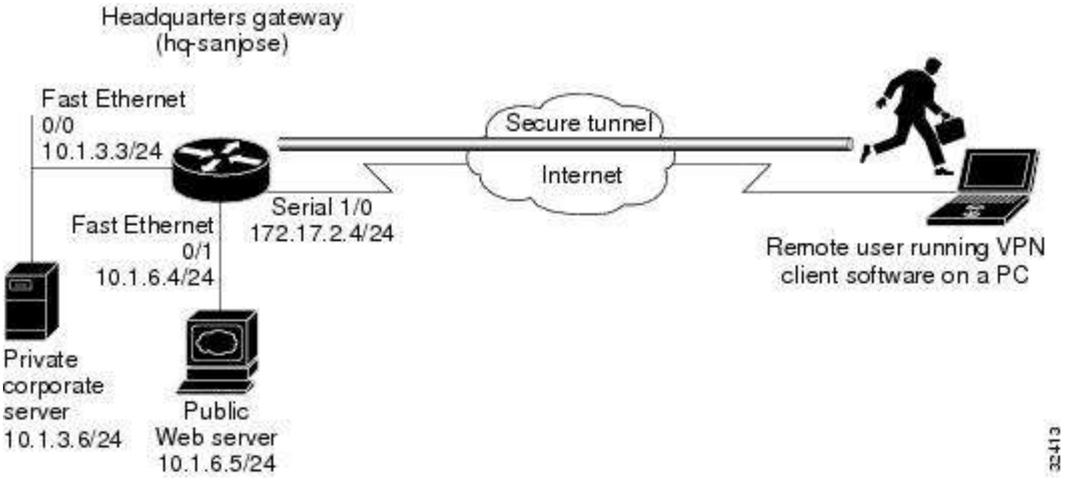

The configuration steps in the following sections are for the headquarters router. Comprehensive configuration examples for the headquarters router are provided in the ["Comprehensive Configuration Examples" section](#). Table 4-1 lists the physical elements of the scenario.

Table 4-1 Physical Elements

Headquarters Network			Remote User		
Site Hardware	**WAN IP Address**	**Ethernet IP Address**	**Site Hardware**	**WAN IP Address**	**Ethernet IP Address**
hq-sanjose	Serial interface 1/0: 172.17.2.4	Fast Ethernet Interface 0/0: 10.1.3.3	PC running VPN client software	Dynamically assigned	—

	255.255.255.0	255.255.255.0 Fast Ethernet Interface 0/1: 10.1.6.4 255.255.255.0			
Corporate server	—	10.1.3.6	—	—	—
Web server	—	10.1.6.5			

Configuring a Cisco IOS VPN Gateway for Use with Cisco Secure VPN Client Software

Using Cisco Secure VPN Client software, a remote user can access the corporate headquarters network through a secure IPSec tunnel. Although Cisco IOS VPN gateways support Cisco Secure VPN Client software, this guide does not explain how to configure your gateway for use with it. For detailed information on configuring client-initiated VPNs using Cisco Secure VPN Client software, refer to the Cisco Secure VPN Client Solutions Guide publication.

Configuring a Cisco IOS VPN Gateway for Use with Microsoft Dial-Up Networking

Using Microsoft Dial-Up Networking (DUN), available with Microsoft Windows 95, Microsoft Windows 98, Microsoft Windows NT 4.0, and Microsoft Windows 2000, a remote user can use Point-to-Point Tunneling Protocol (PPTP) with Microsoft Point-to-Point Encryption (MPPE) to access the corporate headquarters network through a secure tunnel.

Employing PPTP/MPPE, users can use any Internet service provider (ISP) account and any Internet-routable IP address to access the edge of the enterprise network. At the edge, the IP packet is detunneled and the IP address space of the enterprise is used for traversing the internal network. MPPE provides an encryption service that protects the datastream as it traverses the Internet. MPPE is available in two strengths: 40-bit encryption, which is widely available throughout the world, and 128-bit encryption, which may be subject to certain export controls when used outside the United States.

Note PPTP/MPPE is built into Windows DUN1.2 and above. However, 128-bit encryption and stateless (historyless) MPPE is only supported in Windows DUN1.3 or later versions. PPTP/MPPE only supports Cisco Express Forwarding (CEF) and process switching. Regular fast switching is not supported.

Alternatively, a remote user with client software bundled into Microsoft Windows 2000 can use Layer 2 Tunneling Protocol (L2TP) with IPSec to access the corporate headquarters network through a secure tunnel.

Because L2TP is a standard protocol, enterprises can enjoy a wide range of service offerings available from multiple vendors. L2TP implementation is a solution that provides a flexible, scalable remote network access environment without compromising corporate security or endangering mission-critical applications.

Note L2TP is only supported in Microsoft Windows 2000.

This section includes the following topics:

- Configuring PPTP/MPPE
- Verifying PPTP/MPPE
- Configuring L2TP/IPSec

Configuring PPTP/MPPE

PPTP is a network protocol that enables the secure transfer of data from a remote client to a private enterprise server by creating a VPN across TCP/IP-based data networks. PPTP supports on-demand, multiprotocol, virtual private networking over public networks, such as the Internet.

MPPE is an encryption technology developed by Microsoft to encrypt point-to-point links. These PPP connections can be over a dialup line or over a VPN tunnel. MPPE works as a subfeature of Microsoft Point-to-Point Compression (MPPC).

MPPE uses the RC4 algorithm with either 40- or 128-bit keys. All keys are derived from the cleartext authentication password of the user. RC4 is stream cipher; therefore, the sizes of the encrypted and decrypted frames are the same size as the original frame. The Cisco implementation of MPPE is fully interoperable with that of Microsoft and uses all available options, including historyless mode. Historyless mode can increase throughput in high-loss environments such as VPNs.

Note The VAM, available on Cisco 7200 series routers, does not support MPPE.

Note Windows clients must use Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) authentication for MPPE to work. If you are performing mutual authentication with MS-CHAP and MPPE, both sides of the tunnel must use the same password.

This section contains basic steps to configure PPTP/MPPE and includes the following tasks:

- Configuring a Virtual Template for Dial-In Sessions
- Configuring PPTP
- Configuring MPPE

Configuring a Virtual Template for Dial-In Sessions

Using virtual templates, you can populate virtual-access interfaces with predefined customized configurations. To configure your Cisco IOS VPN gateway to create virtual-access interfaces from a virtual template for incoming PPTP calls, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`hq-sanjose(config)# ` **`interface virtual-template`** *`number`*	Creates the virtual template that is used to clone virtual-access interfaces.
Step 2	`hq-sanjose(config-if)# ` **`ip unnumbered`** *`interface-type number`*	Specifies the IP address of the interface the virtual-access interfaces uses.
Step 3	`hq-sanjose(config-if)# ` **`ppp authentication ms-chap`**	Enables MS-CHAP authentication using the local username database. All windows clients using MPPE need to use MS-CHAP.
Step 4	`hq-sanjose(config-if)# ` **`ip local pool default`** *`first-ip-address last-ip-address`*	Configures the default local pool of IP addresses that will be used by clients.
Step 5	`hq-sanjose(config-if)# ` **`peer default ip address pool`** `{`**`default`**`|`**`name`**`}`	Returns an IP address from the default pool to the client.
Step 6	`hq-sanjose(config-if)# ` **`ip mroute-cache`**	Disables fast switching of IP multicast.
Step 7	`hq-sanjose(config-if)# ` **`ppp encrypt mppe`** `{`**`auto`** `|` **`40`** `|` **`128`**`} [`**`passive`** `|` **`required`**`] [`**`stateful`**`]`	(Optional) Enables MPPE encryption on the virtual template[1] if you are using an ISA with Cisco 7200 series router, see the "Configuring MPPE" section. **Note** The VAM, available on Cisco 7200 series routers, does not support MPPE.

[1] Stateful MPPE encryption changes the key every 255 packets. Stateless (historyless) MPPE encryption generates a new key for every packet. Stateless MPPE is only supported in recent versions of Dial-Up Networking (DUN1.3).

Configuring PPTP

To configure a Cisco 7200 series router to accept tunneled PPP connections from a client, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	hq-sanjose(config)# **vpdn-enable**	Enables virtual private dialup networking on the router.
Step 2	hq-sanjose(config)# **vpdn-group 1**	Creates VPDN group 1.
Step 3	hq-sanjose(config-vpdn)# **accept dialin**	Enables the tunnel server to accept dial-in requests.
Step 4	hq-sanjose(config-vpdn-acc-in)# **protocol pptp**	Specifies that the tunneling protocol will be PPTP.
Step 5	hq-sanjose(config-vpdn-acc-in)# **virtual-template** *template-number*	Specifies the number of the virtual template that will be used to clone the virtual-access interface.
Step 6	hq-sanjose(config-vpdn-acc-in)# **exit** hq-sanjose(config-vpdn)# **local name** *localname*	(Optional) Specifies that the tunnel server will identify itself with this local name. If no local name is specified, the tunnel server will identify itself with its host name.

Configuring MPPE

Note The VPN Acceleration Module (VAM) card does not support MPPE.

To configure MPPE on your Cisco 7200 series router (with an ISA), use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	hq-sanjose(config)# **controller isa** *slot/port*	Enter controller configuration mode on the ISM card.
Step 2	hq-sanjose(config-controller)# **encryption mppe**	Enables MPPE encryption.

Verifying PPTP/MPPE

After you complete a connection, enter the **show vpdn tunnel** command or the **show vpdn session** command to verify your PPTP and MPPE configuration.The following example contains typical output:

```
hq-sanjose# show vpdn tunnel | show vpdn session
PPTP Tunnel Information (Total tunnels=1 sessions=1)

LocID RemID Remote Name     State    Remote Address   Port
Sessions
22    22    172.16.230.29  estabd   172.16.230.29    1374  1
```

Configuring L2TP/IPSec

L2TP is an extension of the Point-to-Point (PPP) Protocol and is often a fundamental building block for VPNs. L2TP merges the best features of two other tunneling protocols: Layer 2 Forwarding (L2F) from Cisco Systems and PPTP from Microsoft. L2TP is an Internet Engineering Task Force (IETF) emerging standard.

Note For information on IPSec, see the "Step 3—Configuring Encryption and IPSec" section on page 3-13.

This section contains basic steps to configure L2TP/IPSec and includes the following tasks:

- Configuring a Virtual Template for Dial-In Sessions
- Configuring L2TP
- Configuring Encryption and IPSec

Configuring a Virtual Template for Dial-In Sessions

To configure your Cisco 7200 series router to create virtual-access interfaces from a virtual template for incoming L2TP calls, refer to the "Configuring a Virtual Template for Dial-In Sessions" section.

Note When configuring a virtual template for use with L2TP/IPSec, do not enable MPPE.

Configuring L2TP

To configure a Cisco 7200 series router to accept tunneled L2TP connections from a client, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	hq-sanjose(config)# **vpdn-enable**	Enables virtual private dialup networking on the router.

Step 2	hq-sanjose(config)# **vpdn-group 1**	Creates VPDN group 1.
Step 3	hq-sanjose(config-vpdn)# **accept dialin**	Enables the tunnel server to accept dial-in requests.
Step 4	hq-sanjose(config-vpdn-acc-in)# **protocol l2tp**	Specifies that the tunneling protocol will be L2TP.
Step 5	hq-sanjose(config-vpdn-acc-in)# **virtual-template** *template-number*	Specifies the number of the virtual template that will be used to clone the virtual-access interface.
Step 6	hq-sanjose(config-vpdn-acc-in)# **exit** hq-sanjose(config-vpdn)# **local name** *localname*	(Optional) Specifies that the tunnel server will identify itself with this local name. If no local name is specified, the tunnel server will identify itself with its host name.

Verifying L2TP

Enter the **show vpdn tunnel** command to verify your LT2P configuration.

```
hq-sanjose# show vpdn tunnel
L2TP Tunnel and Session Information (Total tunnels=5
sessions=5)

LocID RemID Remote Name   State  Remote Address  Port
Sessions
  10    8     7206b        est     10.0.0.1       1701       1

LocID RemID TunID Intf    Username    State  Last Chg
Fastswitch
   4     6    10   Vi1      las         est   01:44:39
enabled
```

Configuring Encryption and IPSec

For detailed information on configuring encryption and IPSec, refer to the following sections of this guide:

- Configuring IKE Policies, page 3-15
- Verifying IKE Policies, page 3-19
- Creating Crypto Access Lists, page 3-22
- Verifying Crypto Access Lists, page 3-22
- Defining Transform Sets and Configuring IPSec Tunnel Mode, page 3-23
- Verifying Transform Sets and IPSec Tunnel Mode, page 3-24

Note When using IPSec with L2TP, do not configure IPSec tunnel mode.

- Creating Crypto Map Entries, page 3-25
- Verifying Crypto Map Entries, page 3-26
- Applying Crypto Maps to Interfaces, page 3-27
- Verifying Crypto Map Interface Associations, page 3-28

Note Although the configuration instructions in the listed sections refer to the "Extranet Scenario" section on page 3-4, the same configuration instructions apply to the remote access scenario described in the "Scenario Description" section.

Configuring Cisco IOS Firewall Authentication Proxy

Using the Cisco IOS firewall authentication proxy feature, network administrators can apply specific security policies on a per-user basis. Users can be identified and authorized on the basis of their per-user policy, and access privileges tailored on an individual basis are possible, in contrast with general policy applied across multiple users.

With the authentication proxy feature, users can log into the network or access the Internet via HTTP, and their specific access profiles are automatically retrieved and applied from an authentication server. The user profiles are active only when there is active traffic from the authenticated users.

The authentication proxy is compatible with Network Address Translation (NAT), Context-based Access Control (CBAC), IP Security (IPSec) encryption, and VPN client software.

This section contains basic steps to configure the Cisco IOS Firewall Authentication Proxy and includes the following tasks:

- Configuring Authentication, Authorization, and Accounting
- Configuring the HTTP Server
- Configuring the Authentication Proxy
- Verifying the Authentication Proxy

Configuring Authentication, Authorization, and Accounting

You must configure the authentication proxy for Authentication, Authorization, and Accounting (AAA) services. Use the following commands in global configuration mode to enable authorization and to define the authorization methods:

	Command	Purpose
Step 1	hq-sanjose(config)# **aaa new-model**	Enables the AAA functionality on the router.
Step 2	hq-sanjose(config)# **aaa authentication login default** *TACACS+ RADIUS*	Defines the list of authentication methods at login.
Step 3	hq-sanjose(config)# **aaa authorization auth-proxy default** [*method1* [*method2*...]]	Enables authentication proxy for AAA methods.
Step 4	hq-sanjose(config)# **tacacs-server host** *hostname*	Specifies an AAA server. For RADIUS servers, use the **radius server host** command.
Step 5	hq-sanjose(config)# **tacacs-server key** *sting*	Sets the authentication and encryption key for communications between the router and the AAA server. For RADIUS servers use the **radiusserverkey** command.
Step 6	hq-sanjose(config)# **access-list** *access-list-number* **permit** *tcp* **host** *source* **eq** *tacacs* **host** *destination*	Creates an ACL entry to allow the AAA server return traffic to the firewall. The source address is the IP address of the AAA server, and the destination address is the IP address of the router interface where the AAA server resides.

In addition to configuring AAA on the firewall router, the authentication proxy requires a per-user access profile configuration on the AAA server. To support the authentication proxy, configure the AAA authorization service "auth-proxy" on the AAA server as outlined here:

- Define a separate section of authorization for **auth-proxy** to specify the downloadable user profiles. This does not interfere with other types of service, such as EXEC. The following example shows a user profile on a TACACS server:

```
default authorization = permit
key = cisco
user = newuser1 {
login = cleartext cisco
service = auth-proxy
{
priv-lvl=15
proxyacl#1="permit tcp any any eq 26"
```

```
proxyacl#2="permit icmp any host 60.0.0.2"
proxyacl#3="permit tcp any any eq ftp"
proxyacl#4="permit tcp any any eq ftp-data"
proxyacl#5="permit tcp any any eq smtp"
proxyacl#6="permit tcp any any eq telnet"
```

- The only supported attribute in the AAA server user configuration is **proxyacl#***n.* Use the **proxyacl#***n* attribute when configuring the access lists in the profile. The attribute **proxyacl#***n* is for both RADIUS and TACACS+ attribute-value (AV) pairs.

- The privilege level must be set to 15 for all users.

- The access lists in the user profile on the AAA server must have **permit** only access commands.

- Set the source address to **any** in each of the user profile access list entries. The source address in the access lists is replaced with the source address of the host making the authentication proxy request when the user profile is downloaded to the firewall.

- The supported AAA servers are CiscoSecure ACS 2.1.x for Window NT (where x is a number 0 to 12) and CiscoSecure ACS 2.3 for Windows NT, CiscoSecure ACS 2.2.4 for UNIX and CiscoSecure ACS 2.3 for UNIX, TACACS+ server (vF4.02.alpha), Ascend RADIUS server - radius-980618 (required avpair patch), and Livingston RADIUS server (v1.16).

Configuring the HTTP Server

To use the authentication proxy, you must also enable the HTTP server on the firewall and set the HTTP server authentication method to use AAA. Enter the following commands in global configuration mode:

	Command	Purpose
Step 1	`hq-sanjose(config)#` **`ip http server`**	Enables the HTTP server on the router. The authentication proxy uses the HTTP server to communicate with the client for user authentication.
Step 2	`hq-sanjose(config)#` **`ip http authentication aaa`**	Sets the HTTP server authentication method to AAA.
Step 3	`hq-sanjose(config)#` **`ip http access-class`** *`access-list-number`*	Specifies the access list for the HTTP server.

Configuring the Authentication Proxy

To configure the authentication proxy, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`hq-sanjose(config)#` **`ip auth-proxy`**	Sets the global authentication proxy idle timeout value in minutes. If the timeout expires, user authentication entries are removed, along with any associated

	auth-cache-time *min*	dynamic access lists. The default value is 60 minutes.
Step 2	hq-sanjose(config)# **ip auth-proxy auth-proxy-banner**	(Optional) Displays the name of the firewall router on the authentication proxy login page. The banner is disabled by default.
Step 3	hq-sanjose(config)# **ip auth-proxy name** *auth-proxy-name* **http** [**auth-cache-time** *min*] [**list** *std-access-list*]	Creates authentication proxy rules. The rules define how you apply authentication proxy. This command associates connection initiating HTTP protocol traffic with an authentication proxy name. You can associate the named rule with an access control list, providing control over which hosts use the authentication proxy feature. If no standard access list is defined, the named authentication proxy rule intercepts HTTP traffic from all hosts whose connection initiating packets are received at the configured interface. (Optional) The **auth-cache-time** option overrides the global authentication proxy cache timer. This option provides more control over timeout values for a specific authentication proxy rule. If no value is specified, the proxy rule assumes the value set with the **ip auth-proxy auth-cache-time** command. (Optional) The **list** option allows you to apply a standard access list to a named authentication proxy rule. HTTP connections initiated from hosts in the access list are intercepted by the authentication proxy.
Step 4	hq-sanjose(config)# **interface** *type*	Enters interface configuration mode by specifying the interface type on which to apply the authentication proxy.
Step 5	hq-sanjose(config-if)# **ip auth-proxy** *auth-proxy-name*	In interface configuration mode, applies the named authentication proxy rule at the interface. This command enables the authentication proxy rule with that name.

Verifying the Authentication Proxy

To check the current authentication proxy configuration, use the **show ip auth-proxy configuration** command in privileged EXEC mode. In the following example, the global authentication proxy idle timeout value is set to 60 minutes, the named authentication proxy rule is "pxy," and the idle timeout value for this named rule is 1 minute. The display shows that no host list is specified, meaning that all connections initiating HTTP traffic at the interface are subject to the authentication proxy rule:

```
router# show ip auth-proxy configuration
Authentication cache time is 60 minutes
Authentication Proxy Rule Configuration
```

```
Auth-proxy name pxy
http list not specified auth-cache-time 1 minutes
```

To verify that the authentication proxy is successfully configured on the router, ask a user to initiate an HTTP connection through the router. The user must have authentication and authorization configured at the AAA server. If the user authentication is successful, the firewall completes the HTTP connection for the user. If the authentication is unsuccessful, check the access list and the AAA server configurations.

Display the user authentication entries using the **show ip auth-proxy cache** command in privileged EXEC mode. The authentication proxy cache lists the host IP address, the source port number, the timeout value for the authentication proxy, and the state of the connection. If the authentication proxy state is HTTP_ESTAB, the user authentication was successful.

```
router# show ip auth-proxy cache
Authentication Proxy Cache
Client IP 192.168.25.215 Port 57882, timeout 1, state
HTTP_ESTAB
```

Wait for one minute, which is the timeout value for this named rule, and ask the user to try the connection again. After one minute, the user connection is denied because the authentication proxy has removed the user authentication entry and any associated dynamic ACLs. The user is presented with a new authentication login page and must log in again to gain access through the firewall.

Comprehensive Configuration Examples

This section contains PPTP/MPPE, and L2TP/IPSec comprehensive sample configurations for the headquarters Cisco 7200 series router.

PPTP/MPPE Configuration

```
hq-sanjose# show running-config

Current configuration
!
version 12.0
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname mp12
!
no logging console guaranteed
enable password lab
!
username tester41 password 0 lab41
!
ip subnet-zero
no ip domain-lookup
!
vpdn enable
!
```

```
vpdn-group 1
! Default PPTP VPDN group
accept-dialin
  protocol pptp
  virtual-template 1
local name cisco_pns
!
memory check-interval 1
!
controller ISA 5/0
encryption mppe
!
process-max-time 200
!
interface FastEthernet0/0
ip address 10.1.3.3 255.255.255.0
no ip directed-broadcast
duplex auto
speed auto
!
interface FastEthernet0/1
ip address 10.1.6.4 255.255.255.0
no ip directed-broadcast
duplex auto
speed auto
!
interface Serial1/0
no ip address
no ip directed-broadcast
shutdown
framing c-bit
cablelength 10
dsu bandwidth 44210
!
interface Serial1/1
no ip address
no ip directed-broadcast
shutdown
framing c-bit
cablelength 10
dsu bandwidth 44210
!
interface FastEthernet4/0
no ip address
no ip directed-broadcast
shutdown
duplex half
!
interface Virtual-Template1
ip unnumbered FastEthernet0/0
```

```
no ip directed-broadcast
ip mroute-cache
no keepalive
ppp encrypt mppe 40
ppp authentication ms-chap
!
ip classless
ip route 172.29.1.129 255.255.255.255 1.1.1.1
ip route 172.29.63.9 255.255.255.255 1.1.1.1
no ip http server
!
line con 0
exec-timeout 0 0
transport input none
line aux 0
line vty 0 4
login
!
aaa new-model
aaa authentication login default tacacs+ radius
!Set up the aaa new model to use the authentication proxy.
aaa authorization auth-proxy default tacacs+ radius
!Define the AAA servers used by the router
tacacs-server host 172.31.54.143
tacacs-server key cisco
radius-server host 172.31.54.143
radius-server key cisco
!
! Enable the HTTP server on the router:
ip http server
! Set the HTTP server authentication method to AAA:
ip http authentication aaa
!Define standard access list 61 to deny any host.
access-list 61 deny any
! Use ACL 61 to deny connections from any host to the HTTP
server.
ip http access-class 61
!
!set the global authentication proxy timeout value.
ip auth-proxy auth-cache-time 60
!Apply a name to the authentication proxy configuration rule.
ip auth-proxy name HQ_users http
!
! Apply the authentication proxy rule at an interface.
interface e0
ip address 10.1.1.210 255.255.255.0
ip auth-proxy HQ_users
!
end
```

L2TP/IPSec Configuration

```
hq-sanjose# show running-config

Current configuration:
 !
 version 12.0
 service timestamps debug uptime
 service timestamps log uptime
 no service password-encryption
 !
 hostname LNS
 !
 enable password ww
 !
 username LNS password 0 tunnelpass
 username test@cisco.com password 0 cisco
 ip subnet-zero
 !
 vpdn enable
 !
 vpdn-group 1
  accept dialin l2tp virtual-template 1 remote LAC
  local name LNS
 !
 crypto isakmp policy 1
  authentication pre-share
  group 2
  lifetime 3600
 crypto isakmp key cisco address 172.1.1.1
 !
 crypto ipsec transform-set testtrans esp-des
 !
  !
  crypto map l2tpmap 10 ipsec-isakmp
  set peer 172.1.1.1
  set transform-set testtrans
  match address 101
 !
 interface Ethernet 0/0
  ip address 10.1.3.3 255.255.255.0
  no ip directed-broadcast
  no keepalive
 !
 interface Ethernet 0/1
  no ip address
  no ip directed-broadcast
  shutdown
 !
 interface Virtual-Template1
```

```
  ip unnumbered Ethernet0
  no ip directed-broadcast
  no ip route-cache
  peer default ip address pool mypool
  ppp authentication chap
 !
 interface Serial 1/0
  ip address 172.17.2.4 255.255.255.0
  no ip directed-broadcast
  no ip route-cache
  no ip mroute-cache
  no fair-queue
  clockrate 1300000
  crypto map l2tpmap
 !
 interface Serial 0/0
  no ip address
  no ip directed-broadcast
  shutdown
 !
 ip local pool mypool 172.16.3.1 172.20.10.10
 no ip classless
 !
 access-list 101 permit udp host 172.17.2.4 eq 1701 host
172.1.1.1 eq 1701
 !
 line con 0
  exec-timeout 0 0
  transport input none
 line aux 0
 line vty 0 4
  password cisco
  login
 !
 aaa new-model
 aaa authentication login default tacacs+ radius
 !Set up the aaa new model to use the authentication proxy.
 aaa authorization auth-proxy default tacacs+ radius
 !Define the AAA servers used by the router
 tcacs-server host 172.31.54.143
 tacacs-server key cisco
 radius-server host 172.31.54.143
 radius-server key cisco
 !
 ! Enable the HTTP server on the router:
 ip http server
 ! Set the HTTP server authentication method to AAA:
 ip http authentication aaa
 !Define standard access list 61 to deny any host.
 access-list 61 deny any
 ! Use ACL 61 to deny connections from any host to the HTTP
server.
 ip http access-class 61
 !
 !set the global authentication proxy timeout value.
 ip auth-proxy auth-cache-time 60
 !Apply a name to the authentication proxy configuration rule.
 ip auth-proxy name HQ_users http
 !
 ! Apply the authentication proxy rule at an interface.
 interface e0
ip address 10.1.1.210 255.255.255.0
ip auth-proxy HQ_users
 !
 end
```

Virtual private network

From Wikipedia, the free encyclopedia

Jump to: navigation, search

"VPN" redirects here. For other uses, see VPN (disambiguation).

This article has multiple issues. Please help **improve it** or discuss these issues on the **talk page**.

- It may be **confusing or unclear** for some readers. Tagged since November 2009.
- It is in need of attention from an **expert** on the subject. Tagged since November 2009.

Internet VPN

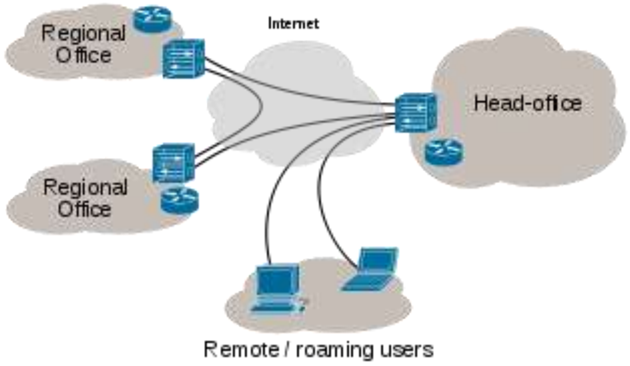

VPN Connectivity overview

A **virtual private network** (**VPN**) is a computer network that uses a public telecommunication infrastructure such as the Internet to provide remote offices or individual users with secure access to their organization's network. It aims to avoid an expensive system of owned or leased lines that can be used by only one organization.

It encapsulates data transfers between two or more networked devices which are not on the same private network so as to keep the transferred data private from other devices on one or more intervening local or wide area networks. There are many different classifications, implementations, and uses for VPNs.

Contents

[hide]

[edit] History

Until the end of the 1990s networked computers were connected through expensive leased lines and/or dial-up phone lines. It could cost thousands of dollars for 56kbps lines or tens of thousands for T1 lines, depending on the distance between the sites.

Virtual Private Networks reduce network costs because they avoid a need for many leased lines that individually connect to the Internet. Users can exchange private data securely, making the expensive leased lines unnecessary.[1]

VPN technologies have a myriad of protocols, terminologies and marketing influences that define them. For example, VPN technologies can differ in:

- The protocols they use to tunnel the traffic
- The tunnel's termination point, i.e., customer edge or network provider edge
- Whether they offer site-to-site or remote access connectivity

- The levels of security provided
- The OSI layer they present to the connecting network, such as Layer 2 circuits or Layer 3 network connectivity

Some classification schemes are discussed in the following sections.

[edit] Security Mechanisms

Secure VPNs use cryptographic tunneling protocols to provide confidentiality by blocking intercepts and packet sniffing, allowing sender authentication to block identity spoofing, and provide message integrity by preventing message alteration.

Secure VPN protocols include the following:

- IPsec (Internet Protocol Security) was originally developed for IPv6, which requires it. This standards-based security protocol is also widely used with IPv4. L2TP frequently runs over IPsec.
- Transport Layer Security (SSL/TLS) can tunnel an entire network's traffic, as it does in the OpenVPN project, or secure an individual connection. A number of vendors provide remote access VPN capabilities through SSL. An SSL VPN can connect from locations where IPsec runs into trouble with Network Address Translation and firewall rules.
- Datagram Transport Layer Security (DTLS), is used in Cisco's next-generation VPN product, Cisco AnyConnect VPN, to solve the issues SSL/TLS has with tunneling over TCP.
- Microsoft's Microsoft Point-to-Point Encryption (MPPE) works with their PPTP and in several compatible implementations on other platforms.
- Microsoft introduced Secure Socket Tunneling Protocol (SSTP) in Windows Server 2008 and Windows Vista Service Pack 1. SSTP tunnels Point-to-Point Protocol (PPP) or L2TP traffic through an SSL 3.0 channel.
- MPVPN (Multi Path Virtual Private Network). Ragula Systems Development Company owns the registered trademark "MPVPN".[2]
- Secure Shell (SSH) VPN -- OpenSSH offers VPN tunneling to secure remote connections to a network or inter-network links. This should not be confused with port forwarding. OpenSSH server provides limited number of concurrent tunnels and the VPN feature itself does not support personal authentication.[3][4][5]

[edit] Authentication

Tunnel endpoints must authenticate before secure VPN tunnels can establish.

User-created remote access VPNs may use passwords, biometrics, two-factor authentication or other cryptographic methods.

Network-to-network tunnels often use passwords or digital certificates, as they permanently store the key to allow the tunnel to establish automatically and without intervention.

[edit] Routing

Tunneling protocols can be used in a point-to-point topology that would theoretically not be considered a VPN, because a VPN by definition is expected to support arbitrary and changing sets of network nodes. But since most router implementations support a software-defined tunnel interface, customer-provisioned VPNs often are simply defined tunnels running conventional routing protocols.

On the other hand provider-provided VPNs (PPVPNs) need to support coexisting multiple VPNs, hidden from one another, but operated by the same service provider.

[edit] PPVPN Building blocks

Depending on whether the PPVPN runs in layer 2 or layer 3, the building blocks described below may be L2 only, L3 only, or combine them both. Multiprotocol Label Switching (MPLS) functionality blurs the L2-L3 identity.

RFC 4026 generalized the following terms to cover L2 and L3 VPNs, but they were introduced in RFC 2547.[6]

Customer edge device. (CE)

a device at the customer premises, that provides access to the PPVPN. Sometimes it's just a demarcation point between provider and customer responsibility. Other providers allow customers to configure it.

Provider edge device (PE)

A PE is a device, or set of devices, at the edge of the provider network, that presents the provider's view of the customer site. PEs are aware of the VPNs that connect through them, and maintain VPN state.

Provider device (P)

A P device operates inside the provider's core network, and does not directly interface to any customer endpoint. It might, for example, provide routing for many provider-operated tunnels that belong to different customers' PPVPNs. While the P device is a key part of implementing PPVPNs, it is not itself VPN-aware and does not maintain VPN state. Its principal role is allowing the service provider to scale its PPVPN offerings, as, for example, by acting as an aggregation point for multiple PEs. P-to-P connections, in such a role, often are high-capacity optical links between major locations of provider.

[edit] User-visible PPVPN services

This section deals with the types of VPN considered in the IETF; some historical names were replaced by these terms.

[edit] OSI Layer 1 services

[edit] Virtual private wire and private line services (VPWS and VPLS)

In both of these services, the provider does not offer a full routed or bridged network, but provides components to build customer-administered networks. VPWS are point-to-point while VPLS can be point-to-multipoint. They can be Layer 1 emulated circuits with no data link structure.

The customer determines the overall customer VPN service, which also can involve routing, bridging, or host network elements.

An unfortunate acronym confusion can occur between Virtual Private Line Service and Virtual Private LAN Service; the context should make it clear whether "VPLS" means the layer 1 virtual private line or the layer 2 virtual private LAN.

[edit] OSI Layer 2 services

Virtual LAN

A Layer 2 technique that allows for the coexistence of multiple LAN broadcast domains, interconnected via trunks using the IEEE 802.1Q trunking protocol. Other trunking protocols have been used but have become obsolete, including Inter-Switch Link (ISL), IEEE 802.10 (originally a security protocol but a subset was introduced for trunking), and ATM LAN Emulation (LANE).

Virtual private LAN service (VPLS)

Developed by IEEE, VLANs allow multiple tagged LANs to share common trunking. VLANs frequently comprise only customer-owned facilities. The former[*clarification needed*] is a layer 1 technology that supports emulation of both point-to-point and point-to-multipoint topologies. The method discussed here extends Layer 2 technologies such as 802.1d and 802.1q LAN trunking to run over transports such as Metro Ethernet.

As used in this context, a VPLS is a Layer 2 PPVPN, rather than a private line, emulating the full functionality of a traditional local area network (LAN). From a user standpoint, a VPLS makes it possible to interconnect several LAN segments over a packet-switched, or optical, provider core; a core transparent to the user, making the remote LAN segments behave as one single LAN.[7]

In a VPLS, the provider network emulates a learning bridge, which optionally may include VLAN service.

Pseudo wire (PW)

PW is similar to VPWS, but it can provide different L2 protocols at both ends. Typically, its interface is a WAN protocol such as Asynchronous Transfer Mode or Frame Relay. In contrast, when aiming to provide the appearance of a LAN contiguous between two or more locations, the Virtual Private LAN service or IPLS would be appropriate.

IP-only LAN-like service (IPLS)

A subset of VPLS, the CE devices must have L3 capabilities; the IPLS presents packets rather than frames. It may support IPv4 or IPv6.

[edit] OSI Layer 3 PPVPN architectures

This section discusses the main architectures for PPVPNs, one where the PE disambiguates duplicate addresses in a single routing instance, and the other, virtual router, in which the PE contains a virtual router instance per VPN. The former approach, and its variants, have gained the most attention.

One of the challenges of PPVPNs involves different customers using the same address space, especially the IPv4 private address space.[8] The provider must be able to disambiguate overlapping addresses in the multiple customers' PPVPNs.

BGP/MPLS PPVPN

In the method defined by RFC 2547, BGP extensions advertise routes in the IPv4 VPN address family, which are of the form of 12-byte strings, beginning with an 8-byte Route Distinguisher (RD) and ending with a 4-byte IPv4 address. RDs disambiguate otherwise duplicate addresses in the same PE.

PEs understand the topology of each VPN, which are interconnected with MPLS tunnels, either directly or via P routers. In MPLS terminology, the P routers are Label Switch Routers without awareness of VPNs.

Virtual router PPVPN

The Virtual Router architecture,[9][10] as opposed to BGP/MPLS techniques, requires no modification to existing routing protocols such as BGP. By the provisioning of logically independent routing domains, the customer operating a VPN is completely responsible for the address space. In the various MPLS tunnels, the different PPVPNs are disambiguated by their label, but do not need routing distinguishers.

Virtual router architectures do not need to disambiguate addresses, because rather than a PE router having awareness of all the PPVPNs, the PE contains multiple virtual router instances, which belong to one and only one VPN.

[edit] Plaintext Tunnels

Main article: Tunneling protocol

Some virtual networks may not use encryption to protect the data contents. While VPNs often provide security, an unencrypted overlay network does not neatly fit within the secure or trusted categorization. For example a tunnel set up between two hosts that used Generic Routing Encapsulation (GRE) would in fact be a virtual private network, but neither secure nor trusted.

Besides the GRE example above, native plaintext tunneling protocols include Layer 2 Tunneling Protocol (L2TP) when it is set up without IPsec and Point-to-Point Tunneling Protocol (PPTP) or Microsoft Point-to-Point Encryption (MPPE).

[edit] Trusted delivery networks

Trusted VPNs do not use cryptographic tunneling, and instead rely on the security of a single provider's network to protect the traffic.

- Multi-Protocol Label Switching (MPLS) is often used to overlay VPNs, often with quality-of-service control over a trusted delivery network.

- Layer 2 Tunneling Protocol (L2TP)[11] which is a standards-based replacement, and a compromise taking the good features from each, for two proprietary VPN protocols: Cisco's Layer 2 Forwarding (L2F)[12] (obsolete as of 2009) and Microsoft's Point-to-Point Tunneling Protocol (PPTP).[13]

From the security standpoint, VPNs either trust the underlying delivery network, or must enforce security with mechanisms in the VPN itself. Unless the trusted delivery network runs among physically secure sites only, both trusted and secure models need an authentication mechanism for users to gain access to the VPN.

[edit] VPNs in mobile environments

Main article: Mobile virtual private network

Mobile VPNs are used in a setting where an endpoint of the VPN is not fixed to a single IP address, but instead roams across various networks such as data networks from cellular carriers or between multiple Wi-Fi access points.[14] Mobile VPNs have been widely used in public safety, where they give law enforcement officers access to mission-critical applications, such as computer-assisted dispatch and criminal databases, as they travel between different subnets of a mobile network.[15] They are also used in field service management and by healthcare organizations,[16] among other industries.

Increasingly, mobile VPNs are being adopted by mobile professionals and white-collar workers who need reliable connections.[16] They allow users to roam seamlessly across networks and in and out of wireless-coverage areas without losing application sessions or dropping the secure VPN session. A conventional VPN cannot survive such events because the network tunnel is disrupted, causing applications to disconnect, time out,[14] or fail, or even cause the computing device itself to crash.[16]

Instead of logically tying the endpoint of the network tunnel to the physical IP address, each tunnel is bound to a permanently associated IP address at the device. The mobile VPN software handles the necessary network authentication and maintains the network sessions in a manner transparent to the application and the user.[14] The Host Identity Protocol (HIP), under study by the Internet Engineering Task Force, is designed to support mobility of hosts by separating the role of IP addresses for host identification from their locator functionality in an IP network. With HIP a mobile host maintains its logical connections established via the host identity identifier while associating with different IP addresses when roaming between access networks.

[edit] See also

- Opportunistic encryption
- Split tunneling
- Mediated VPN
- OpenVPN
- Tinc (protocol)
- SIJOVPN
- SSL VPN

[edit] References

1. ^ Feilner, Markus. "Chapter 1 - VPN—Virtual Private Network". OpenVPN: Building and Integrating Virtual Private Networks: Learn How to Build Secure VPNs Using this Powerful Open Source Application. Packt Publishing.
2. ^ Trademark Applications and Registrations Retrieval (TARR)
3. ^ OpenBSD ssh manual page, VPN section
4. ^ Unix Toolbox section on SSH VPN
5. ^ Ubuntu SSH VPN how-to
6. ^ E. Rosen & Y. Rekhter (March 1999). "RFC 2547 BGP/MPLS VPNs". Internet Engineering Task Forc (IETF). http://www.ietf.org/rfc/rfc2547.txt.
7. ^ *Ethernet Bridging (OpenVPN)*, http://openvpn.net/index.php/access-server/howto-openvpn-as/214-how-to-setup-layer-2-ethernet-bridging.html
8. ^ Address Allocation for Private Internets, RFC 1918, Y. Rekhter *et al.*,February 1996
9. ^ RFC 2917, *A Core MPLS IP VPN Architecture*
10. ^ RFC 2918, K. Muthukrishnan & A. Malis (September 2000)
11. ^ Layer Two Tunneling Protocol "L2TP", RFC 2661, W. Townsley *et al.*,August 1999
12. ^ IP Based Virtual Private Networks, RFC 2341, A. Valencia *et al.*, May 1998
13. ^ Point-to-Point Tunneling Protocol (PPTP), RFC 2637, K. Hamzeh *et al.*,July 1999
14. ^ *a b c* Phifer, Lisa. "Mobile VPN: Closing the Gap", *SearchMobileComputing.com*, July 16, 2006.
15. ^ Willett, Andy. "Solving the Computing Challenges of Mobile Officers", *www.officer.com*, May, 2006.
16. ^ *a b c* Cheng, Roger. "Lost Connections", *The Wall Street Journal*, December 11, 2007.

[edit] External links

- JANET UK "Different Flavours of VPN: Technology and Applications"
- Virtual Private Network Consortium - a trade association for VPN vendors
- CShip VPN-Wiki/List

VPN - Virtual private Networking, an overview

- A VPN is a secure, private communication tunnel between `two or more devices across a public network (like the Internet). These VPN devices can be either a computer running VPN software or a special device like a VPN enabled router. It allows your home computer to be connected to your office network or can allow two home computers in different locations to connect to each over the Internet.

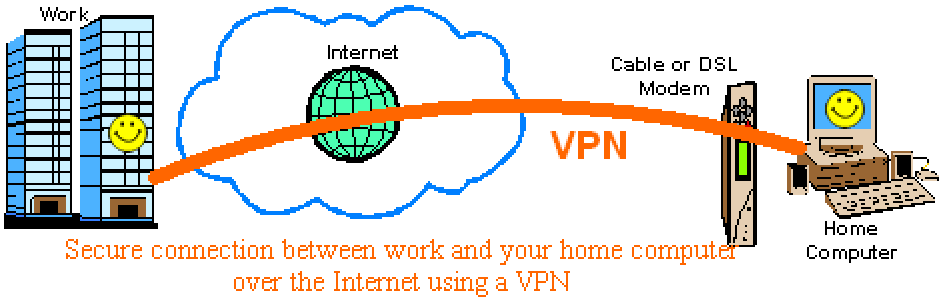

Secure connection between work and your home computer over the Internet using a VPN

- [Network diagrams were made with SmartDraw - click for SmartDraw info.]

- Even though a VPN's data travels across a public network like the Internet, it is secure because of very strong encryption. If anyone 'listens' to the VPN communications, they will not understand it because all the data is encrypted. In addition, VPN's monitor their traffic in very sophisticated ways that ensure packets never get altered while traveling across the public network. Encryption and data verification is very CPU intensive.

- **Clients and Servers**
- A VPN server is a piece of hardware or software that can acts as a gateway into a whole network or a single computer. It is generally 'always on' and listening for VPN clients to connect to it.
- A VPN Client is most often a piece of software but can be hardware too. A client initiates a 'call' to the server and logs on. Then the client computer can server network can communicate. They are on the same 'virtual' network. Many broadband routers can 'pass' one or more VPN sessions from your LAN to the Internet. Each router handles this differently.

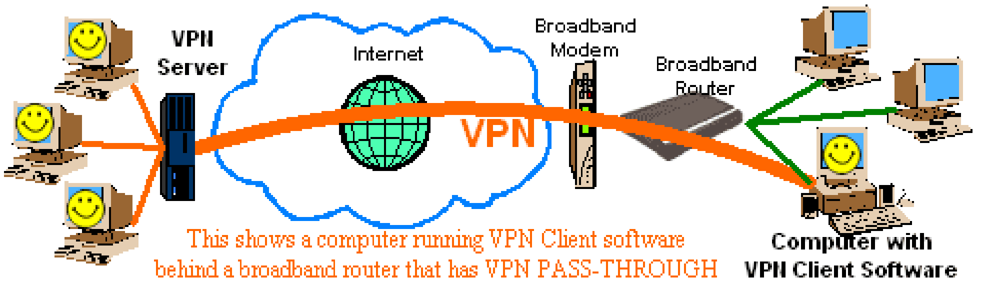

This shows a computer running VPN Client software behind a broadband router that has VPN PASS-THROUGH

- [Network diagrams were made with SmartDraw - click for SmartDraw info.]

- **VPN Software**
- VPN 'server' software is rather rare. Windows Server level operating systems like 'Windows 2000 Server' have a 'VPN server' built in. I know if no software products priced for home or small business that allows you to set up a VPN server.

- VPN ‘client’ software is much more common. When loaded on your computer, this software allows you create a secure VPN tunnel across the Internet and into another network fronted by a VPN server.
-
- **VPN Languages**
- There are two major 'languages' or protocols that VPN's speak. Microsoft uses PPTP or Point to Point Tunneling Protocol and most everyone else uses IPSec - Internet Protocol Security. Most broadband routers can pass PPTP traffic by forwarding port 1723 but IPSec is more complex. If your router does not explicitly support IPSEC pass through, then even placing your computer in the DMZ might not work.
-
- PPTP has 'good' encryption and also features 'authentication' for verifying a user ID and password. IPSec is pureley an encryption model and is mutch safer but does not include authentication routines. A third standard, L2TP is IPSec with authentication built in.
-
- **Broadband Routers with VPN Servers**
- Until recently, VPN server hardware was VERY expensive. As home networks become more sophisticated, the demand for home level VPN’s increase. At the end of 2001, the home network industry responded by adding VPN servers into some broadband routers. These products are often priced at under $300 (us) and some are as inexpensive as $170.

VPN Reviews
2/18/2002: SnapGear Lite+ VPN Router (read)
2/12/2002: ZyWall 1 - Firewall and Router with VPN (read)
1/24/2002: Multitech RouteFinder RF550VPN (read)
1/3/2002: Draytek Vigor2200E & USB VPN Routers (read)

-
- VPN functionality is very processor intensive and most broadband routers have somewhat slow processors in them. Broadband router based VPN servers are often limited in throughput because of their microprocessors. Most have a maximum VPN throughput of around .6Mbps or 600Kbps.
-

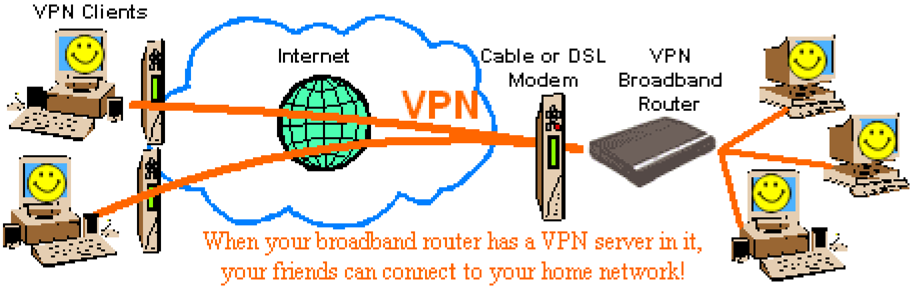

When your broadband router has a VPN server in it, your friends can connect to your home network!

-
- **More info about VPN Routers soon!**
-
-
- **Outside Links for more info**
- VPN Consortium
- Microsofts VPN Pages
- L2TP at Cisco
- VPN Primer at NetGear
- IPSec FAQ at ZyXEL
- VPN Labs Loads of VPN Info
-

HOW TO: Set Up Multiple-Device (Multilink) Dialing in Windows XP

This article describes how to configure multiple-device dialing in Windows XP.

With Windows XP, you can use multiple modems to connect to your Internet service provider (ISP) to increase the total speed of your transfers. Multiple-device dialing (also known as Multilink PPP, modem aggregation, or Multilink) causes multiple physical links to be combined into one logical link. Typically, two or more ISDN lines or modem links are bundled together for greater bandwidth. You might use this feature if you do not have access to DSL or cable services.

Multilink is enabled automatically in Windows XP Home Edition and Windows XP Professional.

Back to the top

Requirements

To use multiple device dialing:

- Your ISP must support synchronization of multiple modems.
- You need to install multiple modems.
- You need a separate phone line for each modem.

Note that a single ISDN adapter can act as multiple devices because ISDN includes two 64-Kbps B channels, which can be used independently or together.

Back to the top

Configuring Multiple-Device Dialing

The Network Connections feature performs Point-to-Point (PPP) Multilink dialing over multiple ISDN, X.25, or modem lines. This feature combines multiple physical links into a logical bundle; the resulting aggregate link increases your connection bandwidth. To dial multiple devices, both your connection and your remote access server must have Multilink enabled.

Network Connections can dynamically control the use of lines that are using Multilink. By allocating lines only as they are required, excess bandwidth is eliminated. You can configure the conditions under which extra lines are dialed, and underused lines are hung up, by changing Network Connections settings.

Note that if you use Multilink to dial a server that requires callback, only one of your Multilink devices is called back. This occurs because you can store only one number in a user account. Therefore, only one device connects and all other devices do not connect; your connection loses Multilink functionality. You can avoid this problem:

- If the phonebook entry for the Multilink connection uses a standard modem configuration, and the remote access server that your connection is calling uses more than one line for the same number.
- If the phonebook entry for the Multilink connection is ISDN with two channels that have the same phone number.

To configure a connection:

1. Click **Start**, click **Control Panel**, click **Network and Internet Connections**, and then click **Network Connections**.
2. Click the connection that you want to configure (for example, a dial-up connection), and then, under **Network Tasks**, click **Change settings of this connection**.
3. Click the **General** tab, and then click each device that you want to use for this connection.
4. Use one or more of the following steps:
   - To configure dialing devices, phone numbers, the host address, country or region codes, or dialing rules, click the **General** tab.
   - To configure dialing and redialing options, or X.25 parameters, click the **Options** tab.
   - To configure identity authentication, data encryption, or terminal window and scripting options, click the **Security** tab.
   - To configure the remote access server and protocols that are used for this connection, click the **Networking** tab. Also, click **Settings** and select the **Negotiate multi-link for single link connections** check box.
   - To enable or disable Internet Connection Sharing, Internet Connection Firewall, and on-demand dialing, click the **Advanced** tab.

Notes:

- Depending on the type of connection that you are configuring, different options and tabs appear in the connection's properties.
- For more information about a specific item on a tab, right-click the item, and then click **What's This?**

Back to the top

Using Routing and Remote Access servers with DHCP

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Using Routing and Remote Access servers with DHCP

The DHCP Server service can be deployed along with the Routing and Remote Access service to provide remote access clients with a dynamically assigned IP address during connection. When these services are used together on the same server computer, the information provided during dynamic configuration is provided differently than in the case of typical DHCP configuration for LAN-based clients.

In LAN environments, DHCP clients negotiate and receive the following configuration information, based entirely on settings configured in the DHCP console for the DHCP server:

- A leased IP address provided from the available address pool of an active scope on the DHCP server. The DHCP server directly manages and distributes the address to the LAN-based DHCP client.
- Additional parameters and other configuration information provided through assigned DHCP options in the address lease. The values and list of options used correspond to option types configured and assigned on the DHCP server.

When a Routing and Remote Access server provides dynamic configuration for dial-up clients, it first performs the following steps:

- When the Routing and Remote Access server starts with the **Use DHCP to assign remote TCP/IP addresses** option, it instructs the DHCP client to obtain 10 IP addresses from a DHCP server.
- The remote access server utilizes the first of these 10 IP addresses obtained from the DHCP server for the remote access server interface.
- The remaining nine addresses are allocated to TCP/IP-based clients as they dial in to establish a session with the remote access server.

IP addresses that are freed when remote access clients disconnect are reused. When all 10 IP addresses are used, the remote access server obtains 10 more from a DHCP server. When the Routing and Remote Access service is stopped, all IP addresses obtained through DHCP are released.

When the Routing and Remote Access server uses this type of proactive caching of DHCP address leases for dial-up clients, it records the following information for each lease response it obtains from the DHCP server:

- The IP address of the DHCP server
- The client leased IP address (for later distribution to the Routing and Remote Access client)
- The time at which the lease was obtained
- The time at which the lease expires
- The duration of the lease

All other DHCP option information returned by the DHCP server (such as server, scope, or reservation options) is discarded. When the client dials in to the server and requests an IP address (that is, when **Server Assigned IP Address** is selected), it uses a cached DHCP lease to provide the dial-up client with dynamic IP address configuration.

When the IP address is provided to the dial-up client, the client is unaware that the IP address has been obtained through this intermediate process between the DHCP server and the Routing and Remote Access server. The Routing and Remote Access server maintains the lease on behalf of the client. Therefore, the only information that the client receives from the DHCP server is the IP address lease.

In dial-up environments, DHCP clients negotiate and receive dynamic configuration using the following modified behavior:

- A leased IP address from the Routing and Remote Access server cache of DHCP scope addresses. The Routing and Remote Access server obtains and renews its cached address pool with the DHCP server.
- If additional parameters and other configuration information provided through assigned DHCP options in the address lease is normally provided by the DHCP server, this information is returned to the Routing and Remote Access client based on TCP/IP properties configured on the Routing and Remote Access server.

Note

- DHCP servers running Windows Server 2003 provide a predefined user class, the **Default Routing and Remote Access Class**, for assigning options that are specific and provided only to Routing and Remote Access clients.

Lesson 6: Managing and Monitoring Remote Access

Managing and monitoring a remote access server can be done with several tools. In this lesson, you learn about remote access logging, accounting, Netsh, Network Monitor, and various resource kit utilities.

After this lesson, you will be able to

- Explain remote access logging
- Describe accounting
- Explain Netsh
- Understand Network Monitor's role in remote access
- List several resource kit utilities to monitor remote access

Estimated lesson time: 30 minutes

Logging User Authentication and Accounting Requests

Internet Authentication Service (IAS) can create log files based on the authentication and accounting requests received from the Network Access Servers (NASs) by collecting these packets in a centralized location. Setting up and using such log files to track authentication information—such as each accept, reject, and automatic account lockout—can help simplify administration of your service. You can set up and use logs to track accounting information—such as logon and logoff records—to help maintain records for billing purposes (see Figure 12.16).

alt="Figure 12.16" src="images/f12cc16.jpg"
Figure 12.16 *Remote Access logging*

When you set up logging, you can specify the following:

- The requests to be logged

- The file format for the logs
- The frequency with which new logs are started
- The location where the logs are to be maintained

You can also select the types of requests received by the IAS server that are to be logged.

Accounting requests include the following:

- Accounting-on requests, which are sent by the NAS to indicate that the NAS is online and ready to accept connections
- Accounting-off requests, which are sent by the NAS to indicate that the NAS is going offline
- Accounting-start requests, which are sent by the NAS (after the user is accepted by the IAS server) to indicate the start of a user session
- Accounting-stop requests, which are sent by the NAS to indicate the end of a user session

Authentication requests include the following:

- Authentication requests, which are sent by the NAS on behalf of the connecting user. These entries in the log contain only incoming attributes.
- Authentication accepts and rejects, which are sent by IAS to the NAS to indicate whether the user should be accepted or rejected. These entries contain only outgoing attributes.
- Periodic status, to obtain interim accounting requests sent by some NASs during sessions.
- Accounting-interim requests, which are sent periodically by the NAS during a user session (if the acct-interim-interval attribute is configured in the remote access profile on the IAS server to support periodic requests).

Initially, it is recommended that you select the first two options and refine your logging methods after you determine which data best matches your needs.

When you set up your servers, specify whether new logs are started daily, weekly, monthly, or when the log reaches a specific size. You can also specify that a single log is maintained continually (regardless of file size), but this is not recommended. The file naming convention for logs is determined by the log period you select. Because changing this option can result in overwriting of existing logs, you should copy logs to a separate file before changing the log period. By default, the log files are located in the %systemroot%\System32\LogFiles folder, but you have the option of specifying a different location.

Log File Records

Attributes are recorded in Unicode Translation Format-8 (UTF-8) encoding in a comma-delimited format. The format of the records in a log file depends on the file format.

- In IAS-formatted log files, each record starts with a fixed-format header, which consists of the NAS IP address, user name, record date, record time, service name, and computer name, which is followed by attribute-value pairs.
- In database-import log files, each record contains attribute values in a consistent sequence, starting with the computer name and are followed by the service name, record date, and record time. An NAS may not use all of the attributes specified in the database-import log format, but the comma-delimited location for each of these predefined attributes is maintained, even for attributes that have no value specified in a record.

Accounting

Routing and Remote Access can be configured to log accounting information in the following locations:

- Locally stored log files when configured for Windows accounting. The information logged and where it is stored are configured from the properties of the Remote Access Logging folder in the Routing and Remote Access snap-in.
- At a RADIUS server when configured for RADIUS accounting. If the RADIUS server is an IAS server, the log files are stored on the IAS server. The information logged and where it is stored

are configured from the properties of the Remote Access Logging folder in the Internet Authentication Service snap-in.

Configuration of the Routing and Remote Access accounting provider is done from the Security tab from the properties of a remote access router in the Routing and Remote Access snap-in, as shown in Figure 12.17, or by using the Netsh tool.

alt="Figure 12.17" src="images/f12cc17.jpg"
Figure 12.17 *Remote Access accounting*

Netsh Command-Line Tool

Netsh is a command-line and scripting tool for Windows 2000 networking components for local or remote computers. Netsh is supplied with Windows 2000. Netsh allows you to save a configuration script in a text file for archival purposes or for configuring other servers.

Netsh is a shell that can support multiple Windows 2000 components through the addition of Netsh helper dynamic-link libraries (DLLs). A Netsh helper DLL extends Netsh functionality by providing additional commands to monitor or configure a specific Windows 2000 networking component. Each Netsh helper DLL provides a context (a group of commands for a specific networking component). Within each context, subcontexts can exist. For example, within the routing context, the subcontexts IP and IPX exist to group IP routing and IPX routing commands together.

For Routing and Remote Access, Netsh has the following contexts:

- **ras.** Use commands in the ras context to configure remote access configuration.
- **aaaa.** Use commands in the aaaa context to configure the AAAA component used by both Routing and Remote Access and IAS.
- **routing.** Use commands in the routing context to configure IP and IPX routing.
- **interface.** Use commands in the interface context to configure demand-dial interfaces.

Network Monitor

Network Monitor enables you to detect and troubleshoot problems on LANs and on WANs, including Routing and Remote Access links. With Network Monitor you can identify network traffic patterns and network problems. For example, you can locate client-to-server connection problems, find a computer that makes a disproportionate number of work requests, capture frames (packets) directly from the network, display and filter the captured frames, and identify unauthorized users on your network.

Resource Kit Utilities

The following are Resource Kit utilities that make the job of managing and monitoring Routing and Remote Access easier.

RASLIST.EXE

The RASLIST.EXE command-line tool displays Routing and Remote Access server announcements from a network. Raslist listens for Routing and Remote Access server announcements on all active network cards in the computer from which it is run. Its output shows which card received the announcement. Raslist is a monitoring tool. It may take a few seconds for the data to begin to appear; data continues to appear until the tool is closed.

RASSRVMON.EXE

By using the RASSRVMON.EXE tool, you can monitor the remote access server activities on your server in greater detail than the standard Windows tools allow. Rassrvmon provides the following monitoring information:

- Server information, such as the time of first call to server, time of most recent call to server, total calls, total bytes passed through server, peak connection count, total connect time, currently connected users, and their connection information.
- Per Port information, which is the time of first call to port, time of most recent call to port, total connections to this port since server started, total bytes passed on this port, total errors on this port, and current port status.
- Summary information, such as statistics kept for each unique user/computer combination since the start of the monitoring, which include total connect time, total bytes transmitted, connection count, average connect time, and total error count.
- Individual connection information, which includes per-connection statistics for each connection: user name/computer name, IP address, connection establishment time, duration, bytes transmitted, error count, and line speed.

To allow for more flexibility, alerts can be set up to run a program of your choice. This gives you the flexibility to send mail, a page, a network popup, or any other action you can automate with an executable file name or a batch script.

RASUSERS.EXE

RASUSERS.EXE lets you list for a domain or a server all user accounts that have been granted permission to dial in to the network via Routing and Remote Access, a feature of Windows 2000 that implements remote access functionality.

TRACEENABLE.EXE

TRACEENABLE.EXE is a graphical user interface-based tool that enables tracing and displays current tracing options. Windows 2000 Routing and Remote Access has an extensive tracing capability that you can use to troubleshoot complex network problems. Tracing records internal component variables, function calls, and interactions. Separate Routing and Remote Access components can be independently enabled to log tracing information to files (file tracing). You must enable the tracing function by changing settings in the Windows 2000 registry using TRACEENABLE.EXE.

Using TraceENABLE.EXE

As each tracing item is selected in the combo box, the values are displayed. Make your changes, and then click Set. This writes your changes to the registry. To get console tracing, you must turn it on for the component and turn it on with the master check box at the top of the Trace Enable window. For example, you would follow these steps to generate a log file for PPP:

1. Select PPP from the drop-down list.
2. Click Enable File Tracing.
3. Click Set.

   Tracing is now enabled for this component. In most cases the log file is created in %windir%\tracing.

Lesson Summary

Managing and monitoring a remote access server is done with several tools. In this lesson, you learned about remote access logging, accounting, Netsh, Network Monitor, and various resource kit utilities.

Today many companies are enjoying the cost savings inherent in allowing some employees to work from home, while those employees benefit from the convenience of telecommuting. In addition, executives, salespeople and others need to connect to the company network when they go on the road, and/or need to access network resources in the evenings or on the weekends from home. All this adds up to a lot of remote access connections to the organization's network. In this article, we will discuss how to prevent remote connections from creating a security nightmare on your network.

- Published: **May 15, 2003**
- Updated: **May 26, 2004**
- Section: **Articles :: Authentication, Access Control & Encryption**
- Author: **Deb Shinder**
- **Printable Version**
- Adjust font size: + −
- Rating: **3.5/5 - 115 Votes**

- 1 ○
- 2 ○
- 3 ○
- 4 ○
- 5 ○

BOOKMARK

Today many companies are enjoying the cost savings inherent in allowing some employees to work from home, while those employees benefit from the convenience of telecommuting. In addition, executives, salespeople and others need to connect to the company network when they go on the road, and/or need to access network resources in the evenings or on the weekends from home.

All this adds up to a lot of remote access connections to the organization's network. These connections may be made over the phone lines by directly dialing into a remote access server on the network, or they may be made by virtual private networking (VPN), using the Internet to "tunnel" into the corporate network. Either way, security is always an issue when you have people connecting to the network from the outside, because you have less control over offsite computers.

There are a number of things you can do, however, to make remote access connections more secure. In this article, we will discuss how to prevent remote connections from creating a security nightmare on your network.

Assessing Remote Connectivity Needs

Your first step in providing for secure remote access is to carefully evaluate employees' need to connect remotely, and grant access on a per-user basis only to those who have a bona fide need to access the network remotely. Always keep in mind that unless restrictions are put in place, a user connecting to the LAN via remote access can do everything that he/she could do from an onsite computer.

In a Windows 2000 domain, you can set user properties to allow or deny remote access, or control access through remote access policies (which we'll discuss in more detail later). To set user properties to allow dial-in or VPN access, configure the properties sheet for each user account, on the Dial-in tab (accessed from the Active Directory Users and Computers MMC), as shown in Figure A.

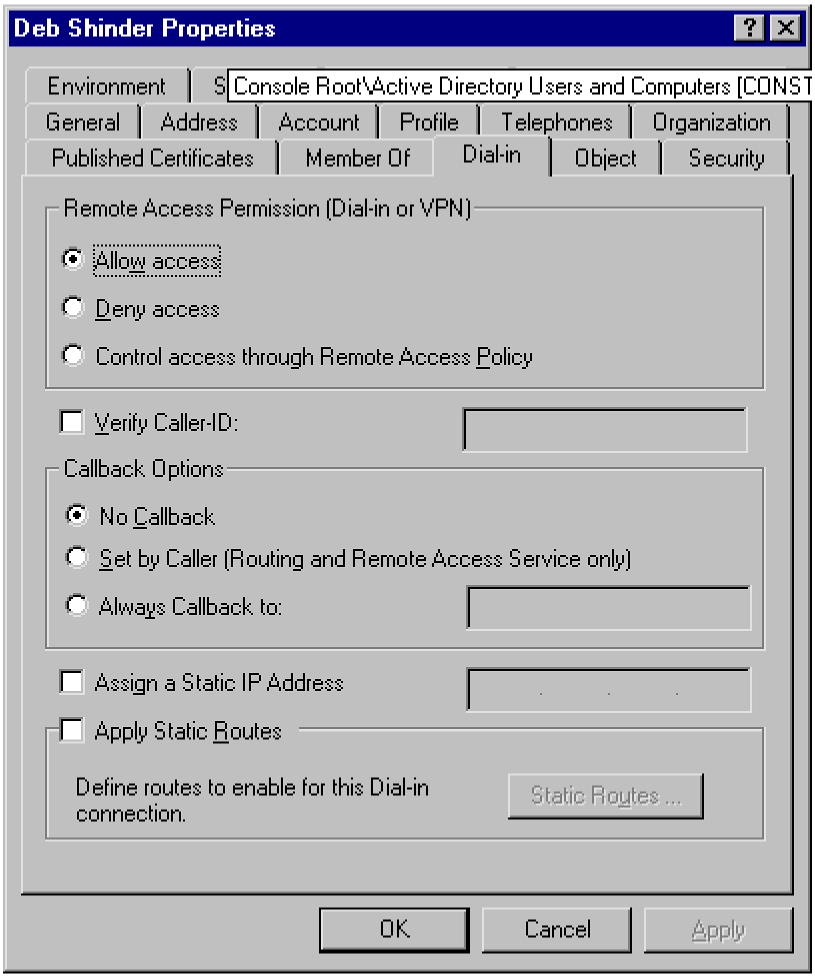

Remote Access Permission (Dial-in or VPN)

● Allow access

○ Deny access

○ Control access through Remote Access Policy

□ Verify Caller-ID:

Callback Options

● No Callback

○ Set by Caller (Routing and Remote Access Service only)

○ Always Callback to:

□ Assign a Static IP Address

□ Apply Static Routes

Define routes to enable for this Dial-in connection.

Figure A: *Setting remote access permissions in user properties*

You can also provide better security for users who work from home or another static location by implementing caller ID verification or setting callback security. To use the former, check the **Verify Caller ID** box and specify a phone number from which the user must dial in. To use the latter, check the **Always Callback to** option button and enter the phone number from which the remote user will connect. The server will hang up and call the user back at that number. Either way, if an unauthorized user manages to discover a legitimate user's password, he/she still won't be able to access the network remotely unless doing so from the legitimate user's location.

Although the two features accomplish the same security objectives, there are a couple of situations in which you would use Callback instead of Caller ID verification: 1) when the caller's or server's phone systems don't support Caller ID and 2) when the remote user is dialing in from a long distance location and you want the server to call back so the company will pay the telephone charges for the session.

After assessing who needs remote access, determine whether you will allow remote users to dial in, connect via VPN, or both. A dialup connection has the security advantage of being a direct connection between the user and the dialup server, so no information is going across the public Internet. VPNs use encryption to protect the confidentiality of information that travels through the public network and provide the "private" aspect. The policies you set in implementing your dialup server or VPN server will determine, to a large extent, the level of security.

For VPN connections, you'll want to consider the protocols your VPN server will support. L2TP tunneling with IPSec encryption is more secure than PPTP (which uses MPPE for encryption); however, not all clients can use L2TP. If all your remote access clients use Windows 2000 or XP (as they should, for best security), your policies can specify L2TP VPN connections only. If you have Windows 9x clients, you may have to allow PPTP connections.

Authentication Considerations

One of your most important security considerations is how remote clients will be authenticated. Authentication, of course, involves verifying the identity of the client computer or user. Remote access authentication protocols are not all created equal.

Windows supports a variety of remote access authentication protocols, ranging from PAP (password authentication protocol), which transmits passwords in plain text and is not secure, to sophisticated authentication methods such as EAP-RADIUS, which relies on a separate authentication server, or EAP-TLS, requiring that the user provide a smart card with a digital certificate.

PAP is disabled by default on the Windows 2000 remote access server, and for best security, you should use only strong authentication. If you use MS-CHAP, use version 2, and set password length and complexity policies to force the use of strong passwords.

Windows 2000/2003 also supports use of security hosts for remote access. This is a device that sits between the remote access client and the remote access server, and provides supplemental authentication (in addition to that of the RAS server). You may have to edit the modem.inf file on the server to link the security host to the server's modem.

Windows Remote Access Security Policies

Unlike NT, which used only the remote access permissions in the user properties to control access, Windows 2000 and 2003 take both the user properties and remote access policies into account. The policies allow you granular control. You can grant remote access only during certain times of the day, or certain days of the week. You can grant VPN access but deny dial-in (or vice versa). You can limit the duration of each remote access session, allow connections only with specified authentication methods, and so forth.

Remote access policies can be set on the remote access server, or the policies for multiple dial-in and VPN servers can be managed through an IAS (Internet Authentication Service) server. To set policies on the remote access server, you use the RRAS MMC (accessed from Administrative Tools), as shown in Figure B.

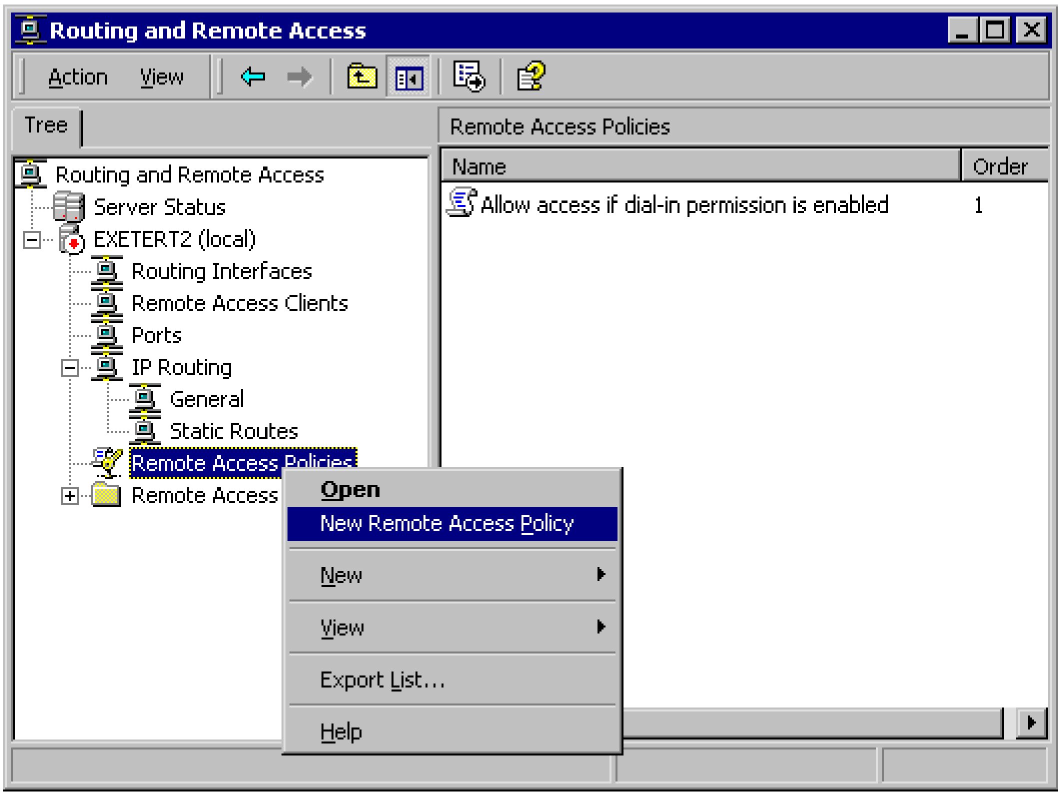

Figure B: *Creating new remote access policies using the RRAS console*

Using Encryption to Secure Dialup Remote Access Connections

The security of remote access connections can be increased by encrypting the data that flows across the phone lines. There are a couple of ways to accomplish this with Windows:

- Use Microsoft Point to Point Encryption (MPPE) with MS-CHAP or EAP-TLS authentication. This is called link encryption, because the data is encrypted only between the routers (gateways) connecting the two networks.
- Use IPSec to encrypt the data all the way from the sending computer to the destination computer. This is called end-to-end encryption.
  To encrypt the data and configure the encryption settings, select the remote access policy in the RRAS console tree, right click it and select **Properties**. In the properties box, click the **Edit Profile button** at the bottom, as shown in Figure C.

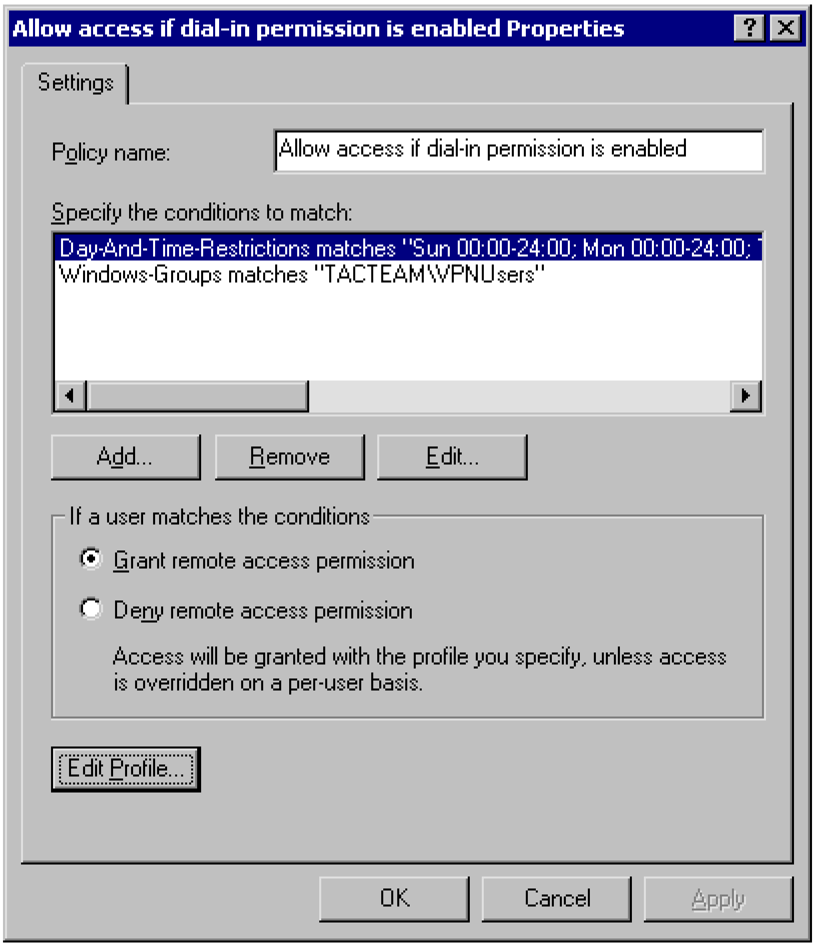

Allow access if dial-in permission is enabled Properties

Settings

Policy name: Allow access if dial-in permission is enabled

Specify the conditions to match:

```
Day-And-Time-Restrictions matches "Sun 00:00-24:00; Mon 00:00-24:00;
Windows-Groups matches "TACTEAM\VPNUsers"
```

If a user matches the conditions

● Grant remote access permission

○ Deny remote access permission

Access will be granted with the profile you specify, unless access is overridden on a per-user basis.

Figure C: *Edit the remote access policy profile to enable data encryption*

In the Edit Profile dialog box, click the **Encryption** tab and check the encryption levels to be allowed by the profile, as shown in Figure D.

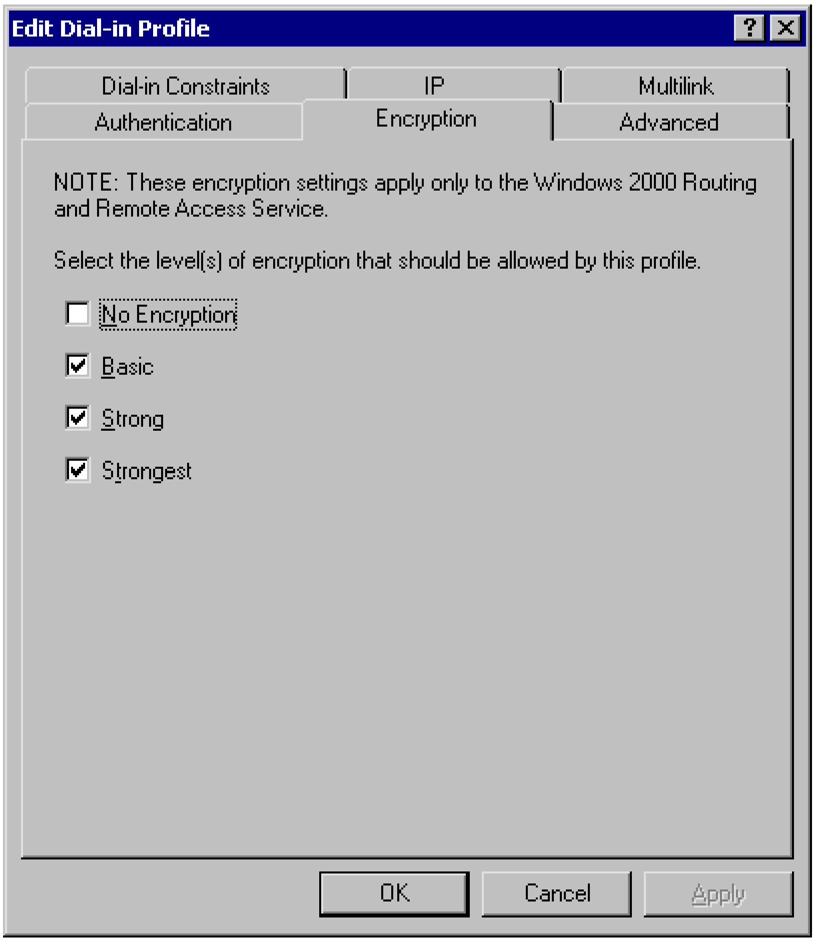

NOTE: These encryption settings apply only to the Windows 2000 Routing and Remote Access Service.

Select the level(s) of encryption that should be allowed by this profile.

□ No Encryption

☑ Basic

☑ Strong

☑ Strongest

Figure D: *Encryption levels are set via the Encryption tab when editing the Dial-in profile*

By clearing the **No Encryption** checkbox, you can require that connections be encrypted.

You can also set the allowed authentication methods for the policy by clicking the **Authentication** tab. This is where you can configure the policy to allow only smart card or certificate-based authentication, for example. At a minimum, you should ensure that the checkboxes labeled **Unencrypted Authentication (PAP, SPAP)** and **Unauthenticated Access** are unchecked, as shown in Figure E.

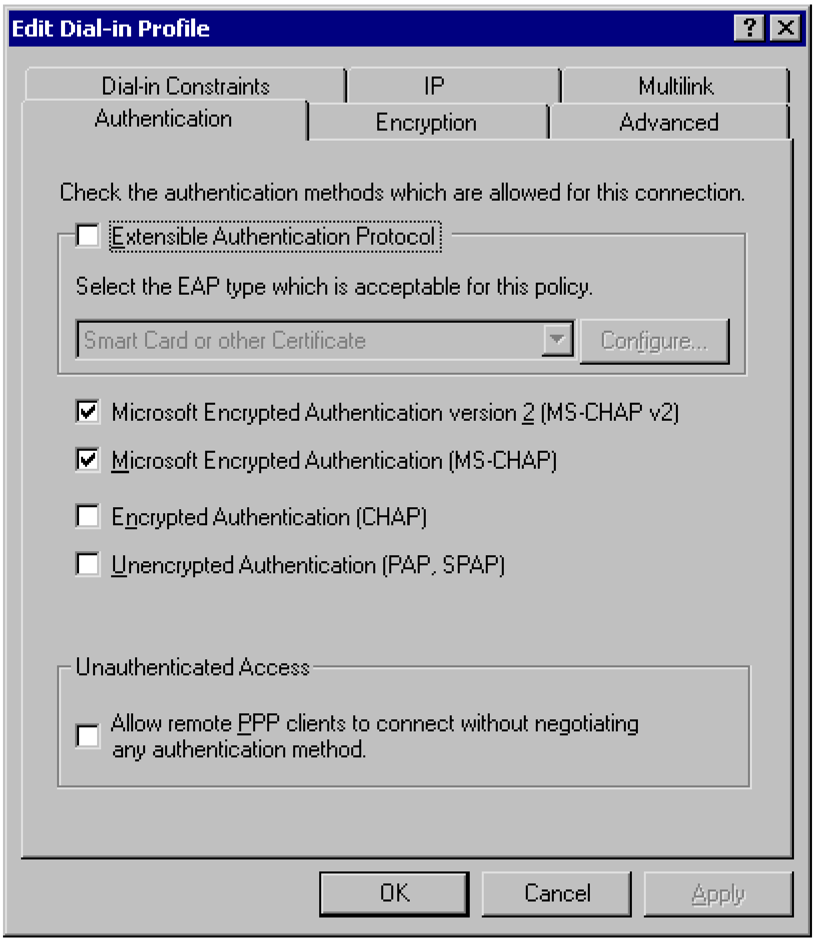

Check the authentication methods which are allowed for this connection.

□ Extensible Authentication Protocol

Select the EAP type which is acceptable for this policy.

Smart Card or other Certificate

☑ Microsoft Encrypted Authentication version 2 (MS-CHAP v2)

☑ Microsoft Encrypted Authentication (MS-CHAP)

□ Encrypted Authentication (CHAP)

□ Unencrypted Authentication (PAP, SPAP)

Unauthenticated Access

□ Allow remote PPP clients to connect without negotiating any authentication method.

Figure E: *Set the allowed authentication methods to disallow unencrypted authentication and unauthenticated access*

Restrictions on dial-in connections can be set by clicking the **Dial-in Constraints** tab. Here you can set idle time limitations, limit the maximum session time, or define the days and times when remote access is allowed, as shown in Figure F.

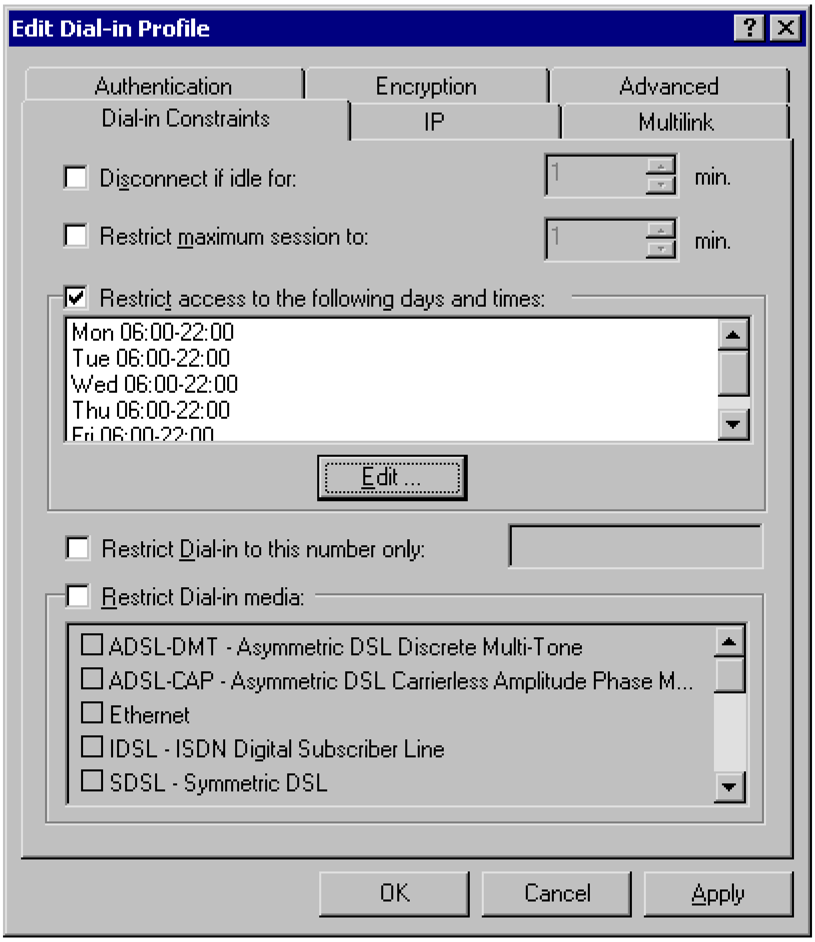

☐ Disconnect if idle for: 1 min.

☐ Restrict maximum session to: 1 min.

☑ Restrict access to the following days and times:

Mon 06:00-22:00
Tue 06:00-22:00
Wed 06:00-22:00
Thu 06:00-22:00
Fri 06:00-22:00

☐ Restrict Dial-in to this number only:

☐ Restrict Dial-in media:

☐ ADSL-DMT - Asymmetric DSL Discrete Multi-Tone
☐ ADSL-CAP - Asymmetric DSL Carrierless Amplitude Phase M...
☐ Ethernet
☐ IDSL - ISDN Digital Subscriber Line
☐ SDSL - Symmetric DSL

Figure F: *You can set a variety of constraints on Dial-in connections*

You can also restrict dial-in connections to a particular phone number and/or restrict the dial-in media type (for instance, to ISDN, modem, or VPN).

Network Protocols encompassing

Troubleshooting the Cisco 200

Table of Contents

Troubleshooting the Cisco 200

Troubleshooting the Cisco 200

This appendix provides general guidelines for troubleshooting the Cisco 200. It also contains a troubleshooting worksheet that lists the information you should have available when seeking technical assistance. This appendix contains the following sections:

Service and Support

For information about obtaining technical assistance with your Cisco 200 product, refer to the *Cisco Information Packet,* which shipped with your product.

Troubleshooting Checklist

The following list of guidelines will help ensure that the Cisco 200 will operate properly:

- When you run INSTALL or SETUP, do not enter any non-numeric characters in the telephone number entry display.

- If you use the CONNECT.BAT file, enter the name of the connecting router at the end of the line. The router name is case-sensitive. For example, if the router name is "HEADQUARTERS" in all capital letters, you must enter:

```
CONNECT HEADQUARTERS
```

not:

```
CONNECT Headquarters
```

- If you manually load the Cisco 200, the name of the connecting router must appear at the end of the ISDN200 statement. The router name is case-sensitive. For example:

```
ISDN200 HEADQUARTERS
```

not:

```
ISDN200 Headquarters
```

- The node MAC address must be unique on the network. Check with your network administrator.
- Select the correct ISDN protocol (for example, NI1).
- The ISDN Adapter must be firmly seated and properly connected, and the jumpers must be set correctly.
- Use a properly configured NT1 if required by your ISDN provider.
- Do not use a terminal adapter; the ISDN Adapter has a built-in terminal adapter.
- Select the proper ISDN line speed (56 kbps or 64 kbps). Verify that the line speed is the same as the line speed of the connecting router.
- When configuring the connecting router, if you are using a 56 kbps ISDN line, include the SERVICE=DATA56 statement in your router's configuration file.
- If you are using CHAP authentication, make sure you are using the correct password. Remember, passwords are case-sensitive.
- Enter the correct SPID, directory name, or both, if required for your ISDN protocol.
- All phone numbers must be correct.
- Ensure that the network protocols you want to use are enabled. Check the ISDN200.CFG file to make sure there is a semicolon (;) in the first position of the line for all the protocols you want to use.
- Ensure that a dialer map is present and configured properly on the connecting router.

Blank Screen and Other Hardware Problems

If the screen is blank, or if your workstation seems to be executing commands slowly, the Cisco 200 software with VLM or NETX loaded might be trying to determine whether the command you entered is a NetWare or a DOS command. Therefore, it will bring up a suspended ISDN connection before executing a command. Be sure you unload VLM or NETX when you finish using the services of the remote file server.

Cannot Connect to Router

Ensure that there are no duplicate network addresses. Addresses are usually assigned by the network administrator.

Troubleshooting the Connecting Router Configuration

The following example illustrates a connecting router configuration for a Cisco 2503 router using PPP multilink with PPP callback and the DSS1 protocol. Have your system administrator use this example if you are having problems connecting using your Cisco 200 and suspect a router configuration problem may be the cause. The lines in bold-face type relate to your installation and are explained in detail in the text following the example.

```
!

version 11.0
```

```
!
hostname HEADQUARTERS
!
username MyHomeOffice password Secret
ipx routing 0000.0c07.ad01
isdn switch-type basic-net3
isdn tei-negotiation first-call
!
interface Ethernet0
ip address 220.6.7.50 255.255.255.0
ipx network 50
!
interface BRI0
ip address 202.6.7.50 255.255.255.0
encapsulation ppp
ipx network 22
no ipx route-cache
ipx watchdog-spoof
dialer enable-timeout 2
dialer map ip 202.6.7.30 name MyHomeOffice CONNECT5551111
dialer map ipx 22.HEADQUARTERS name MyHomeOffice CONNECT5551111
dialer-group 1
isdn caller 5551111
no fair-queue
ppp callback accept
ppp multilink
ppp authentication chap
```

```
Headquarters

map-class dialer CallbackUsers

dialer callback-server ISDN200

!

router rip

network 202.6.7.0

network 210.6.7.0

network 220.6.7.0

!

dialer-list 1 list 100
```

- **hostname** *HEADQUARTERS*
  - "HEADQUARTERS" is the name of the connecting router.
- **username** *MyHomeOffice* password *Secret*
  - "MyHomeOffice" is the name for this Cisco 200.
  - "Secret" is the password for this Cisco 200. This password will be encrypted in the router configuration after the router configuration is saved.
- **dialer map ip 202.6.7.30 name** *MyHomeOffic*e *HEADQUARTERS5551111*
  - "MyHomeOffice" is the name for this Cisco 200.
  - "5551111" is the telephone number of the Cisco 200.
- **dialer map ipx 22.ISDN200** name *MyHomeOffice Headquarters5551111*
  - "0004.5509.1234" is the node MAC address.
  - "MyHomeOffice" is the name for this Cisco 200.
  - "5551111" is the telephone number of the Cisco 200.

Using NETX.EXE with Windows or Windows for Workgroups

If you are using NETX.EXE with Windows or Windows for Workgroups, you need the following files:

- LSL.COM (supplied with Cisco 200 software and installed into the Cisco 200 directory)
- IPXODI.COM
- NETX.EXE

These files should be provided by your system administrator.

Installation Summaries for the Various Platforms

Time Saver This section presents a quick overview of the steps to install the Cisco 200 software for use with the various operating systems and NetWare platforms.

DOS and NetWare ODI Drivers (NETX)

Step 1 Install the Cisco 200 software.

Step 2 Copy the LSL.COM, IPXODI.COM, and NETX.EXE files to the directory in which you installed the Cisco 200 software.

DOS and NetWare VLM Drivers

Step 1 Install the NetWare Client software version 1.2 or later in the designated directory.

Step 2 Rename NET.CFG to NET.OLD.

Step 3 Install the Cisco 200 software in the same directory in which you installed the NetWare Client software.

Windows and NetWare ODI Drivers (NETX)

Step 1 Install the Cisco 200 software in a separate directory.

Step 2 Copy the LSL.COM, IPXODI.COM, and NETX.EXE files to the same directory in which you installed the Cisco 200 software.

Windows and NetWare VLM Drivers

Step 1 Install the NetWare Client software version 1.2 or later in a separate directory.

Step 2 Rename NET.CFG to NET.OLD.

Step 3 Install the Cisco 200 software into the same directory in which you installed the NetWare Client software.

Windows for Workgroups and NetWare ODI Drivers (NETX)

Step 1 Install the Cisco 200 software in a separate directory.

Step 2 Install or copy the LSL.COM, IPXODI.COM, and NETX.EXE files to the same directory in which you installed the Cisco 200 software.

Windows for Workgroups and NetWare VLM Drivers

Step 1 Install the NetWare Client software version 1.2 or later in a separate directory.

Step 2 Rename NET.CFG to NET.OLD.

Step 3 Install the Cisco 200 software in the same directory in which you installed the NetWare Client software.

Troubleshooting Worksheet

Before you request technical assistance, complete the following Troubleshooting Worksheet. Support personnel will ask you to provide this information when you first contact them, and may ask you to fax this information. Please print clearly.

Table C-1: Troubleshooting Worksheet for the Cisco 200

Company Data	
Company name	
Contact	
Phone, fax, or Internet address	
PC	
Manufacturer/model	
Type (386, etc.)/speed (MHz)	
Total amount of memory	
Operating system and version (DOS 6.2, Windows 3.1, etc.)	
ISDN Adapter	
I/O port address, IRQ	
Serial number	
Connecting Router(s)	
Model	
Software version	
NT1 (North America)	
Manufacturer/model	
Miscellaneous	
Network software type and version	

Messages	
Are there any messages appearing on the console?	Yes No
If so, what is the exact wording of the message(s)?	
Network Diagram	
On a separate sheet of paper, draw a diagram showing the configuration of your Cisco 200 network. Include the following:	
• Network addresses	
• Telephone numbers	
• Contents (preferably printouts) of the CONFIG.SYS, AUTOEXEC.BAT, ISDN200.CFG, and NET.CFG files, and any files referred to by a CALL statement.	

Basic Network Application Troubleshooting With Wireshark (Ethereal)

Network protocol analysis is a technique used to view, in real time, the raw data sent and received over a network interface. This is useful for troubleshooting network configuration and network application problems. It is also useful when developing new network protocols. In this article, I'll go over the very basics of troubleshooting a network application with Wireshark (Ethereal).

Wireshark Overview

Ethereal is a open source protocol analyzer initially written by Gerald Combs. It was renamed Wireshark in 2006 because of trademarks'copyrights held on the name Ethereal. Wireshark is now maintained and enhanced by hundreds of people worldwide.

What can it do?

Several years ago I worked on a compliance project to archive all instant messaging (IM) traffic flowing through the company network for the Yahoo! IM, MSN and AOL IM networks. I was a member of the Messaging and Groupware team which was responsible for (among other things) email, IM and other collaborative software systems.

M&G, as the team was called, was not responsible for the network, dns servers and certainly not the firewall. We had to communicate our intentions and coordinate our efforts with the groups responsible for these other systems in order to complete the project. Here is a short (read incomplete) list of what needed to happen - or so we thought.

1. DNS servers needed to report the address of our IM archive server. The archive server works like a web proxy server and will open connections to the real IM server on behalf of the clients.
2. The firewall needed to allow the IM archive server to open connections to the internet on specific ports to a specific server.
3. IM clients needed to be reconfigured to "connect directly to the internet" as opposed to using a socks proxy or http tunneling. The IM client would send a query to resolve a DNS name of an IM server (for example login.oscar.aol.com) to an IP address. The company's internal DNS servers would return the address of our IM proxy which would then open a connection to the real IM server out on the internet. And our proxy would log every conversation detail along the way. Big Brother would be alive and well.

After all of the details were worked out and implemented, it was time to test. But when we flipped the switch, it did not work! I queried our DNS server to check that the proper DNS A records were created, and they were. I checked that I could connect to the IM proxy server on the ports used by each of my IM clients, and I could. I then checked to see if the proxy server could connect to the internet IM servers on the required ports, and I could not.

So obviously the problem was with the firewall, right? I called up the firewall team and complained to them about how they had not opened up the ports I'd requested. However, they swore they had opened up all reports to the specified servers I'd requested. So what was happening?

Well, that's when the product support guy for IM Manager (the IM proxy) suggested we use Ethereal to "see" what was going on. To shorten a long story, with Ethereal we were able to determine that the proxy server could not open a connection to the public IM server since the firewall rules in place only allowed it to open connection to the internet for a specific host but on any port. The hosts we'd given our firewall team were really a sort of traffic cop. It would instruct the IM client to connect to another machine. Since this host was not in our firewall rule set, the connection was not allowed. In the end we requested that the proxy servers be allowed to open connections to any host on the required ports. The required changes were made to the firewall rule sets and everything worked as expected.

This is a classic example of troubleshooting with Wireshark.

What can’t it do?

Just as with any tool, Wireshark can be used for some things and not others. Here is a list of some of the things Wireshark cannot do:

1. It cannot be used to map out a network. Take a look at the NMAP tool for that functionality.
2. It does not generate network data – it is a passive tool. Tools like NMAP, ping, and traceroute are examples of tools that generate network data. These tools are active.
3. It can only show detailed information about protocols it actually understands. The good news is that it understands a great many protocols. It is also extensible, so you can add protocol support for ones it doesn’t understand. Otherwise you will only be able to see a hexdump of data it has captured.
4. It can only capture data as well as the OS'Interface'Interface driver supports. An example of this is capturing data over wireless networks. This does not work well (or at all) for some software and hardware combinations.

Installing Wireshark

Wireshark is an open source application and may be downloaded for free from www.wireshark.org. Installation is straight forward. To install on Windows using the executable package:

1. Double click the installer file.

2. Click the ‘Next’ button at the Welcome screen.

3. Click the ‘I Agree’ button to accept the licensing terms.
4. Click the ‘Next’ button to accept the defaults at the Choose Components dialog box.

5. Click the ‘Next’ button at the Select Additional Tasks dialog box.

6. Click the ‘Next’ button at the Choose Install Location dialog box.

7. At this point, the installer will ask if you want to install WinPcap. Ensure that the Install Winpcap checkbox is selected and click the ‘Next’ button.
8. The Wireshark installation will now begin copying files to your system.
9. The WinPcap installer will launch during Wireshark installation. Click the ‘Next’ button at the Welcome screen.

10. Click the ‘Next’ button at the WinPcap Setup Wizard screen.
11. Click the ‘I Agree’ at the License Agreement screen.

12. Click the ‘Finish’ button to close the WinPcap installer.
13. Click the ‘Next’ button on the Wireshark Installation Complete dialog box.
14. Click the ‘Finish’ button to close the Wireshark installer.

Running Wireshark on Windows

Launching Wireshark

Running Wireshark on Windows is a simple matter of double clicking the shortcut on the start menu. This will open the Wireshark main screen.

WHAT WILL YOU DO WITH VISUAL STUDIO 2010?
LEARN MORE
REPLAY
Microsoft Visual Studio 2010

Wireshark Interface

The Wireshark interface is fairly simple considering what it can do.

1. **Title bar** – this will contain different information depending on what Wireshark is doing. If it is capturing network data, it will show the interface that is in use. If it is displaying data from a previous capture, the name of the file containing the captured data will be shown (untitled is shown if a capture was performed, stopped and not saved). Otherwise it will show the application name: **Wireshark Network Protocol Analyzer**
2. **Menu bar** – Menu bar providing access to application features
    a. *File* – Functions for working with captured data such as saving and exporting to different file formats
    b. *Edit* – Functions for finding packets, setting the time reference, and setting preferences
    c. *View* - Functions for modifying how Wireshark displays information such as which windows are open
    d. *Go* – Functions for navigating to specific packets
    e. *Capture* – Functions for starting and stopping captures, saving filters and working with network interfaces
    f. *Analyze* – Functions for interpreting and filtering captured data
    g. *Statistics* – Functions to statistically analyze captured data
    h. *Help* – access to product help
3. **Main tool bar** – Shortcuts to frequently used functions in the menu bar
4. **Filter tool bar** – Quick access to filter functions
5. **Packet list pane** – Displays all the packets in the current capture file.
6. **Packet details pane** – Shows a more detailed view of the packet currently selected in the Packet List pane
7. **Packet bytes pane** – A hexdump view of the packet currently selected in the Packet List pane

8. **Status bar** – Provides informational messages and feedback to the user

Sample Wireshark Capture

In this example, I will start a Wireshark capture on my wired laptop interface. I will then launch Thunderbird to retrieve email from Comcast and GMail

1. First launch **Wireshark**.
2. Then select Capture->Interfaces from the menu bar.
3. 3) This will bring up the Interfaces dialog box. Select the interface you want to use. This is important since Wireshark (as with any protocol analyzer) can only capture data from a network it is physically connected to. I will be using the wired Ethernet adapter in my laptop so I will choose the Intel adapter in the list. Click the 'Start' button. Capturing will now begin. After a short while, you will see the main Wireshark window (the packet list, detail and byte panes) fill with data.
4. Now I will launch Thunderbird and login to both my GMail and Comcast mail accounts. At this point I will wait for all my mail to download and then I will stop the network capture by selecting Capture->Stop from the menu bar. Click File and Save to save this capture to disk after all data is captured. 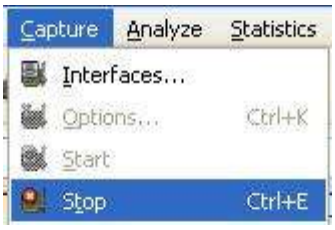

5. I have just captured two complete pop3 sessions with Wireshark. To single out the pop session information I will apply a filter. In the filter bar enter the following text and press the 'apply' button: tcp.port eq 110. This will limit the display to traffic on tcp port 110 (the POP port). Also notice that Wireshark "understands" the Post Office Protocol, so it will interpret bits of information such as POP commands and even authentication information. I do not connect to the Comcast mail server using SSL so my password is contained in the trace in clear text. I had to choose this screenshot wisely! I've actually used this to troubleshoot end user client connection problem to pop and imap servers.
6. Scrolling through the filtered captured data only shows a conversation between two hosts; my laptop and the Comcast mail server. What happened to gmail? Well, I use SSL with my GMail account and SSL POP connections are associated with port 995 not 110. In the filter bar enter the following text and press the 'apply' button: tcp.port eq 995. This will show all the POP over SSL traffic. But notice that no other details are available about the application protocol. The protocols in use on port 995 are TCP, SSL and TLS. You will see some packets

dealing with key exchange, but that is all to do with the security negotiations associated with SSL'TLS. All the Application data is encrypted.

Closing remarks

This is only the tip of the iceberg. I do not have a need to use this tool very often, but when I do need it, it is there and it can be a life saver – or a job saver.

For more information on this tool visit the Wireshark website at www.wireshark.org. The documentation on this site is quite extensive and should get you up to speed fairly quickly.

Also check out the documentation for whatever protocol you are sniffing. For the internet protocols, the RFC's are a must read. You can find these at the Internet Engineering Task Force website at www.ietf.org.

Finally, if you are sniffing an undocumented protocol like nrpc (used by Lotus Notes'Domino) you may want to spend time researching a similar protocol that is documented. This may help you to understand what you are seeing in Wireshark related to the undocumented protocol.

Related Articles

- Portscanning with NMAP
- Quickly Find Local Open Ports
- Quickly Find Local Open Ports - GUI
- Quickly Find Remote Open Ports
- Quickly Find Remote Open Ports - GUI

Recent Networking Forum threads

Got a question? Post it on our Windows Networking Forums!

Related Articles

- 5 Critical VMware ESX CLI Network Troubleshooting Commands
- Home Network Setup – What are the possible configuration settings for a home/SOHO network with 3-4 computers and an ADSL Internet connection?
- Recovery and Troubleshooting Options in Windows XP
- Troubleshooting Dcpromo Errors

Transmission Control Protocol

From Wikipedia, the free encyclopedia
Jump to: navigation, search

The **Transmission Control Protocol** (**TCP**) is one of the core protocols of the Internet Protocol Suite. TCP is one of the two original components of the suite, complementing the Internet Protocol (IP), and therefore the entire suite is commonly referred to as *TCP/IP*. TCP provides the service of exchanging data directly between two network hosts, whereas IP handles addressing and routing message across one or more networks. In particular, TCP provides reliable, ordered delivery of a stream of bytes from a program on one computer to another program on another computer. TCP is the protocol that major Internet applications rely on, applications such as the World Wide Web, e-mail, and file transfer. Other applications, which do not require reliable data stream service, may use the User Datagram Protocol (UDP) which provides a datagram service that emphasizes reduced latency over reliability.

Internet Protocol Suite

Application Layer

BGP · DHCP · DNS · FTP · HTTP · IMAP ·
IRC · LDAP · MGCP · NNTP · NTP · POP ·
RIP · RPC · RTP · SIP · SMTP · SNMP · SSH ·
Telnet · TLS/SSL · XMPP ·

(more)

Transport Layer

TCP · UDP · DCCP · SCTP · RSVP · ECN ·

(more)

Internet Layer

IP (IPv4, IPv6) · ICMP · ICMPv6 · IGMP · IPsec ·

(more)

Link Layer

ARP/InARP · NDP · OSPF · Tunnels (L2TP) ·
PPP · Media Access Control (Ethernet, DSL,
ISDN, FDDI) · (more)

This box: view · talk · edit

Contents

[hide]

[edit] Historical origin

In May, 1974, the Institute of Electrical and Electronic Engineers (IEEE) published a paper entitled "*A Protocol for Packet Network Interconnection.*"[1] The paper's authors, Vinton G. Cerf and Bob Kahn, described an internetworking protocol for sharing resources using packet-switching among the nodes. A central control component of this model was the *Transmission Control Program* that incorporated both connection-oriented links and datagram services between hosts. The monolithic Transmission Control Program was later divided into a modular architecture consisting of the *Transmission Control Protocol* at the connection-oriented layer and the *Internet Protocol* at the internetworking (datagram) layer. The model became known informally as *TCP/IP*, although formally it was henceforth called the *Internet Protocol Suite*.

[edit] Network function

TCP provides a communication service at an intermediate level between an application program and the Internet Protocol (IP). That is, when an application program desires to send a large chunk of data across the Internet using IP, instead of breaking the data into IP-sized pieces and issuing a series of IP requests, the software can issue a single request to TCP and let TCP handle the IP details.

IP works by exchanging pieces of information called packets. A packet is a sequence of octets and consists of a *header* followed by a *body*. The header describes the packet's destination and, optionally, the routers to use for forwarding until it arrives at its destination. The body contains the data IP is transmitting.

Due to network congestion, traffic load balancing, or other unpredictable network behavior, IP packets can be lost, duplicated, or delivered out of order. TCP detects these problems, requests retransmission of lost packets, rearranges out-of-order packets, and even helps minimize network congestion to reduce the occurrence of the other problems. Once the TCP receiver has finally reassembled a perfect copy of the data originally transmitted, it passes that datagram to the application program. Thus, TCP abstracts the application's communication from the underlying networking details.

TCP is utilized extensively by many of the Internet's most popular applications, including the World Wide Web (WWW), E-mail, File Transfer Protocol, Secure Shell, peer-to-peer file sharing, and some streaming media applications.

TCP is optimized for accurate delivery rather than timely delivery, and therefore, TCP sometimes incurs relatively long delays (in the order of seconds) while waiting for out-of-order messages or retransmissions of lost messages. It is not particularly suitable for real-time applications such as Voice over IP. For such applications, protocols like the Real-time Transport Protocol (RTP) running over the User Datagram Protocol (UDP) are usually recommended instead.[2]

TCP is a reliable stream delivery service that guarantees delivery of a data stream sent from one host to another without duplication or losing data. Since packet transfer is not reliable, a technique known as positive acknowledgment with retransmission is used to guarantee reliability of packet transfers. This fundamental technique requires the receiver to respond with an acknowledgment message as it receives the data. The sender keeps a record of each packet it sends, and waits for acknowledgment before sending the next packet. The sender also keeps a timer from when the packet was sent, and retransmits a packet if the timer expires. The timer is needed in case a packet gets lost or corrupted.[2]

TCP consists of a set of rules: for the protocol, that are used with the Internet Protocol, and for the IP, to send data "in a form of message units" between computers over the Internet. At the same time that IP takes care of handling the actual delivery of the data, TCP takes care of keeping track of the individual units of data transmission, called *segments*, that a message is divided into for efficient routing through the network. For example, when an HTML file is sent from a Web server, the TCP software layer of that server divides the sequence of octets of the file into segments and forwards them individually to the IP software layer (Internet Layer). The Internet Layer encapsulates each TCP segment into an IP packet by adding a header that includes (among other data) the destination IP address. Even though every packet has the same destination address, they can be routed on different paths through the network. When the client program on the destination computer receives them, the TCP layer (Transport Layer) reassembles the individual segments and ensures they are correctly ordered and error free as it streams them to an application.

[edit] TCP segment structure

A TCP segment consists of a segment *header* and a *data* section. The TCP header contains 10 mandatory fields, and an optional extension field (*Options*, pink background in table).

The data section follows the header. Its contents are the payload data carried for the application. The length of the data section is not specified in the TCP segment header. It can be calculated by subtracting the combined length of the TCP header and the encapsulating IP segment header from the total IP segment length (specified in the IP segment header).

TCP Header

Bit offset	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	Source port																Destination port															
32	Sequence number																															
64	Acknowledgment number																															
96	Data offset				Reserved				CWR	ECE	URG	ACK	PSH	RST	SYN	FIN	Window Size															
128	Checksum																Urgent pointer															
160 ...	Options (if Data Offset > 5) ...																															

- Source port (16 bits) – identifies the sending port
- Destination port (16 bits) – identifies the receiving port
- Sequence number (32 bits) – has a dual role:
    - If the `SYN` flag is set, then this is the initial sequence number. The sequence number of the actual first data byte (and the acknowledged number in the corresponding ACK) are then this sequence number plus 1.
    - If the `SYN` flag is clear, then this is the accumulated sequence number of the first data byte of this packet for the current session.

- Acknowledgment number (32 bits) – if the ACK flag is set then the value of this field is the next sequence number that the receiver is expecting. This acknowledges receipt of all prior bytes (if any). The first ACK sent by each end acknowledges the other end's initial sequence number itself, but no data.
- Data offset (4 bits) – specifies the size of the TCP header in 32-bit words. The minimum size header is 5 words and the maximum is 15 words thus giving the minimum size of 20 bytes and maximum of 60 bytes, allowing for up to 40 bytes of options in the header. This field gets its name from the fact that it is also the offset from the start of the TCP segment to the actual data.
- Reserved (4 bits) – for future use and should be set to zero
- Flags (8 bits) (aka Control bits) – contains 8 1-bit flags

    - CWR (1 bit) – Congestion Window Reduced (CWR) flag is set by the sending host to indicate that it received a TCP segment with the ECE flag set and had responded in congestion control mechanism (added to header by RFC 3168).
    - ECE (1 bit) – ECN-Echo indicates

    - If the SYN flag is set, that the TCP peer is ECN capable.
    - If the SYN flag is clear, that a packet with Congestion Experienced flag in IP header set is received during normal transmission (added to header by RFC 3168).

    - URG (1 bit) – indicates that the Urgent pointer field is significant
    - ACK (1 bit) – indicates that the Acknowledgment field is significant. All packets after the initial SYN packet sent by the client should have this flag set.
    - PSH (1 bit) – Push function. Asks to push the buffered data to the receiving application.
    - RST (1 bit) – Reset the connection
    - SYN (1 bit) – Synchronize sequence numbers. Only the first packet sent from each end should have this flag set. Some other flags change meaning based on this flag, and some are only valid for when it is set, and others when it is clear.
    - FIN (1 bit) – No more data from sender

- Window (16 bits) – the size of the *receive window*, which specifies the number of bytes (beyond the sequence number in the acknowledgment field) that the receiver is currently willing to receive (*see Flow control and Window Scaling*)
- Checksum (16 bits) – The 16-bit checksum field is used for error-checking of the header and data
- Urgent pointer (16 bits) – if the URG flag is set, then this 16-bit field is an offset from the sequence number indicating the last urgent data byte
- Options (Variable 0-320 bits, divisible by 32) – The length of this field is determined by the data offset field. Options 0 and 1 are a single byte (8 bits) in length. The remaining options indicate the total length of the option (expressed in bytes) in the second byte.[3] Some options may only be sent when SYN is set; they are indicated below as [SYN].

    - **0** (8 bits) - End of options list

- **1** (8 bits) - No operation (NOP, Padding) This may be used to align option fields on 32-bit boundaries for better performance.
- **2**,**4**,*SS* (32 bits) - Maximum segment size (*see maximum segment size*) `[SYN]`
- **3**,**3**,*S* (24 bits) - Window scale (*see window scaling for details*) `[SYN]`[4]
- **4**,**2** (16 bits) - Selective Acknowledgement permitted. `[SYN]` (*See selective acknowledgments for details*)[5]
- **5**,*N*,*BBBB*,*EEEE*,... (variable bits, *N* is either 10, 18, 26, or 34)- Selective ACKnowlegement (SACK)[6] These first two bytes are followed by a list of 1-4 blocks being selectively acknowledged, specified as 32-bit begin/end pointers.
- **8**,**10**,*TTTT*,*EEEE* (80 bits)- Timestamp and echo of previous timestamp (*see TCP Timestamps for details*)[7]
- **14**,**3**,*S* (24 bits) - TCP Alternate Checksum Request. `[SYN]`[8]
- **15**,*N*,... (variable bits) - TCP Alternate Checksum Data.

(The remaining options are obsolete, experimental, not yet standardized, or unassigned)

[edit] Protocol operation

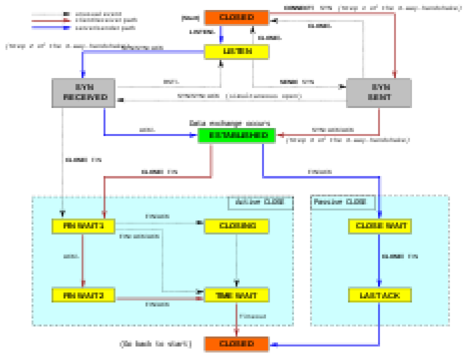

A Simplified TCP State Diagram. See TCP EFSM diagram for a more detailed state diagram including the states inside the ESTABLISHED state.

TCP protocol operations may be divided into three phases. Connections must be properly established in a multi-step handshake process (*connection establishment*) before entering the *data transfer* phase. After data transmission is completed, the *connection termination* closes established virtual circuits and releases all allocated resources.

A TCP connection is managed by an operating system through a programming interface that represents the local end-point for communications, the *Internet socket*. During the lifetime of a TCP connection it undergoes a series of state changes:

1. LISTEN : In case of a server, waiting for a connection request from any remote client.
2. SYN-SENT : waiting for the remote peer to send back a TCP segment with the SYN and ACK flags set. (usually set by TCP clients)

3. SYN-RECEIVED : waiting for the remote peer to send back an acknowledgment after having sent back a connection acknowledgment to the remote peer. (usually set by TCP servers)
4. ESTABLISHED : the port is ready to receive/send data from/to the remote peer.
5. FIN-WAIT-1
6. FIN-WAIT-2
7. CLOSE-WAIT
8. CLOSING
9. LAST-ACK
10. TIME-WAIT : represents waiting for enough time to pass to be sure the remote peer received the acknowledgment of its connection termination request. According to RFC 793 a connection can stay in TIME-WAIT for a maximum of four minutes.
11. CLOSED

[edit] Connection establishment

To establish a connection, TCP uses a three-way handshake. Before a client attempts to connect with a server, the server must first bind to a port to open it up for connections: this is called a passive open. Once the passive open is established, a client may initiate an active open. To establish a connection, the three-way (or 3-step) handshake occurs:

1. The active open is performed by the client sending a SYN to the server. It sets the segment's sequence number to a random value A.
2. In response, the server replies with a SYN-ACK. The acknowledgment number is set to one more than the received sequence number (A + 1), and the sequence number that the server chooses for the packet is another random number, B.
3. Finally, the client sends an ACK back to the server. The sequence number is set to the received acknowledgement value i.e. A + 1, and the acknowledgement number is set to one more than the received sequence number i.e. B + 1.

At this point, both the client and server have received an acknowledgment of the connection.

[edit] Resource usage

Most implementations allocate an entry in a table that maps a session to a running operating system process. Because TCP packets do not include a session identifier, both endpoints identifies the session using the client's address and port. Whenever a packet is received, the TCP implementation must perform a lookup on this table to find the destination process.

The number of sessions in the server side is limited only by memory and can grow as new connections arrive, but the client must allocate a random port before sending the first SYN to the server. This port remains allocated during the whole conversation, and effectively limits the number of outgoing connections from each of the client's IP addresses. If an application fails to properly close unrequired connections, a client can run out of resources and become unable to establish new TCP connections, even from other applications.

Both endpoints must also allocate space for unacknowledged packets and received (but unread) data.

[edit] Data transfer

There are a few key features that set TCP apart from User Datagram Protocol:

- Ordered data transfer - the destination host rearranges according to sequence number[2]
- Retransmission of lost packets - any cumulative stream not acknowledged is retransmitted[2]
- Error-free data transfer (The checksum in UDP is optional)
- Flow control - limits the rate a sender transfers data to guarantee reliable delivery. The receiver continually hints the sender on how much data can be received (controlled by the sliding window). When the receiving host's buffer fills, the next acknowledgment contains a 0 in the window size, to stop transfer and allow the data in the buffer to be processed.[2]
- Congestion control [2]

[edit] Reliable transmission

TCP uses a *sequence number* to identify each byte of data. The sequence number identifies the order of the bytes sent from each computer so that the data can be reconstructed in order, regardless of any fragmentation, disordering, or packet loss that may occur during transmission. For every payload byte transmitted the sequence number must be incremented. In the first two steps of the 3-way handshake, both computers exchange an initial sequence number (ISN). This number can be arbitrary, and should in fact be unpredictable to defend against TCP Sequence Prediction Attacks.

TCP primarily uses a *cumulative acknowledgment* scheme, where the receiver sends an acknowledgment signifying that the receiver has received all data preceding the acknowledged sequence number. Essentially, the first byte in a segment's data field is assigned a sequence number, which is inserted in the sequence number field, and the receiver sends an acknowledgment specifying the sequence number of the next byte they expect to receive. For example, if computer A sends 4 bytes with a sequence number of 100 (conceptually, the four bytes would have a sequence number of 100, 101, 102 and 103 assigned) then the receiver would send back an acknowledgment of 104 since that is the next byte it expects to receive in the next packet.

In addition to cumulative acknowledgments, TCP receivers can also send selective acknowledgments to provide further information (*see selective acknowledgments*).

If the sender infers that data has been lost in the network, it retransmits the data.

[edit] Error detection

Sequence numbers and acknowledgments cover discarding duplicate packets, retransmission of lost packets, and ordered-data transfer. To assure correctness a checksum field is included (*see TCP segment structure for details on checksumming*).

The TCP checksum is a weak check by modern standards. Data Link Layers with high bit error rates may require additional link error correction/detection capabilities. The weak checksum is partially compensated for by the common use of a CRC or better integrity check at layer 2, below both TCP and IP, such as is used in PPP or the Ethernet frame. However, this does not mean that the 16-bit TCP checksum is redundant: remarkably, introduction of errors in packets between CRC-protected hops is common, but the end-to-end 16-bit TCP checksum catches most of these simple errors [9]. This is the end-to-end principle at work.

[edit] Flow control

TCP uses an end-to-end flow control protocol to avoid having the sender send data too fast for the TCP receiver to receive and process it reliably. Having a mechanism for flow control is essential in an environment where machines of diverse network speeds communicate. For example, if a PC sends data to a hand-held PDA that is slowly processing received data, the PDA must regulate data flow so as not to be overwhelmed.[2]

TCP uses a sliding window flow control protocol. In each TCP segment, the receiver specifies in the **receive window** field the amount of additional received data (in bytes) that it is willing to buffer for the connection. The sending host can send only up to that amount of data before it must wait for an acknowledgment and window update from the receiving host.

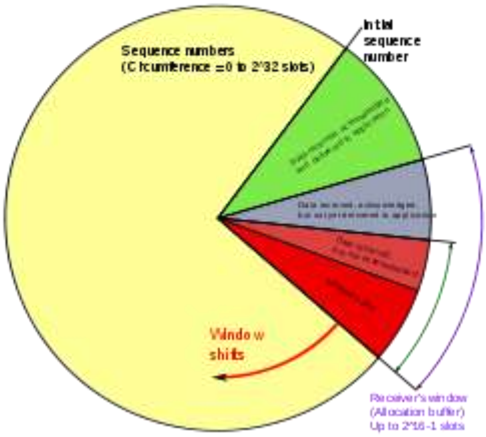
Initial sequence number
Sequence numbers (Circumference = 0 to 2^32 slots)
Window shifts
Receiver's window (Allocation buffer) Up to 2^16-1 slots

TCP sequence numbers and receive windows behave very much like a clock. The receive window shifts each time the receiver receives and acknowledges a new segment of data. Once it runs out of sequence numbers, the sequence number loops back to 0.

When a receiver advertises a window size of 0, the sender stops sending data and starts the **persist timer**. The persist timer is used to protect TCP from a deadlock situation that could arise if a subsequent window size update from the receiver is lost, and the sender cannot send more data until receiving a new window size update from the receiver. When the persist timer expires, the TCP sender attempts recovery by sending a small packet so that the receiver responds by sending another acknowledgement containing the new window size.

If a receiver is processing incoming data in small increments, it may repeatedly advertise a small receive window. This is referred to as the silly window syndrome, since it is inefficient to send only a few bytes of data in a TCP segment, given the relatively large overhead of the TCP header. TCP senders and receivers typically employ flow control logic to specifically avoid repeatedly sending small segments. The sender-side silly window syndrome avoidance logic is referred to as Nagle's algorithm.

[edit] Congestion control

The final main aspect of TCP is congestion control. TCP uses a number of mechanisms to achieve high performance and avoid 'congestion collapse', where network performance can fall by several orders of magnitude. These mechanisms control the rate of data entering the network, keeping the data flow below a rate that would trigger collapse.

Acknowledgments for data sent, or lack of acknowledgments, are used by senders to infer network conditions between the TCP sender and receiver. Coupled with timers, TCP senders and receivers can alter the behavior of the flow of data. This is more generally referred to as congestion control and/or network congestion avoidance.

Modern implementations of TCP contain four intertwined algorithms: Slow-start, congestion avoidance, fast retransmit, and fast recovery (RFC 5681).

In addition, senders employ a **retransmission timeout** (RTO) that is based on the estimated round-trip time (or RTT) between the sender and receiver, as well as the variance in this round trip time. The behavior of this timer is specified in RFC 2988. There are subtleties in the estimation of RTT. For example, senders must be careful when calculating RTT samples for retransmitted packets; typically they use Karn's Algorithm or TCP timestamps (see RFC 1323). These individual RTT samples are then averaged over time to create a Smoothed Round Trip Time (SRTT) using Jacobson's algorithm. This SRTT value is what is finally used as the round-trip time estimate.

Enhancing TCP to reliably handle loss, minimize errors, manage congestion and go fast in very high-speed environments are ongoing areas of research and standards development. As a result, there are a number of TCP congestion avoidance algorithm variations.

[edit] Maximum segment size

The Maximum segment size (MSS) is the largest amount of data, specified in bytes, that TCP is willing to send in a single segment. For best performance, the MSS should be set small enough to avoid IP fragmentation, which can lead to excessive retransmissions if there is packet loss. To try to accomplish this, typically the MSS is negotiated using the MSS option when the TCP connection is established, in which case it is determined by the maximum transmission unit (MTU) size of the data link layer of the networks to which the sender and receiver are directly attached. Furthermore, TCP senders can use Path MTU discovery to infer the minimum MTU along the network path between the sender and receiver, and use this to dynamically adjust the MSS to avoid IP fragmentation within the network.

[edit] Selective acknowledgments

Relying purely on the cumulative acknowledgment scheme employed by the original TCP protocol can lead to inefficiencies when packets are lost. For example, suppose 10,000 bytes are sent in 10 different TCP packets, and the first packet is lost during transmission. In a pure cumulative acknowledgment protocol, the receiver cannot say that it received bytes 1,000 to 9,999 successfully, but failed to receive the first packet, containing bytes 0 to 999. Thus the sender may then have to resend all 10,000 bytes.

To solve this problem TCP employs the *selective acknowledgment (SACK)* option, defined in RFC 2018, which allows the receiver to acknowledge discontinuous blocks of packets that were received correctly, in addition to the sequence number of the last contiguous byte received successively, as in the basic TCP acknowledgment. The acknowledgement can specify a number of *SACK blocks*, where each SACK block is conveyed by the starting and ending sequence numbers of a contiguous range that the receiver correctly received. In the example above, the receiver would send SACK with sequence numbers 1,000 and 9,999. The sender thus retransmits only the first packet, bytes 0 to 999.

An extension to the SACK option is the "duplicate-SACK" option, defined in RFC 2883. An out-of-order packet delivery can often falsely indicate the TCP sender of lost packet and, in turn, the TCP sender retransmits the suspected-to-be-lost packet and slow down the data delivery to prevent network congestion. The TCP sender undoes the action of slow-down, that is a recovery of the original pace of data transmission, upon receiving a D-SACK that indicates the retransmitted packet is duplicate.

The SACK option is not mandatory and it is used only if both parties support it. This is negotiated when connection is established. SACK uses the optional part of the TCP header (*see TCP segment structure for details*). The use of SACK is widespread - all popular TCP stacks support it. Selective acknowledgment is also used in Stream Control Transmission Protocol (SCTP).

[edit] Window scaling

Main article: TCP window scale option

For more efficient use of high bandwidth networks, a larger TCP window size may be used. The TCP window size field controls the flow of data and its value is limited to between 2 and 65,535 bytes.

Since the size field cannot be expanded, a scaling factor is used. The TCP window scale option, as defined in RFC 1323, is an option used to increase the maximum window size from 65,535 bytes to 1 Gigabyte. Scaling up to larger window sizes is a part of what is necessary for TCP Tuning.

The window scale option is used only during the TCP 3-way handshake. The window scale value represents the number of bits to left-shift the 16-bit window size field. The window scale value can be set from 0 (no shift) to 14 for each direction independently. Both sides must send the option in their SYN segments to enable window scaling in either direction.

Some routers and packet firewalls rewrite the window scaling factor during a transmission. This causes sending and receiving sides to assume different TCP window sizes. The result is non-stable traffic that may be very slow. The problem is visible on some sending and receiving sites behind the path of defective routers.[10]

[edit] TCP Timestamps

TCP timestamps, defined in RFC 1323, help TCP compute the round-trip time between the sender and receiver. Timestamp options include a 4-byte timestamp value, where the sender inserts its current value of its timestamp clock, and a 4-byte echo reply timestamp value, where the receiver generally inserts the most recent timestamp value that it has received. The sender uses the echo reply timestamp in an acknowledgment to compute the total elapsed time since the acknowledged segment was sent.[2]

TCP timestamps are also used to help in the case where TCP sequence numbers encounter their 2^{32} bound and "wrap around" the sequence number space. This scheme is known as *Protect Against Wrapped Sequence* numbers, or *PAWS* (see RFC 1323 for details). Furthermore, the Eifel detection algorithm, defined in RFC 3522, which detects unnecessary loss recovery requires TCP timestamps.

[edit] Out of band data

One is able to interrupt or abort the queued stream instead of waiting for the stream to finish. This is done by specifying the data as *urgent*. This tells the receiving program to process it immediately, along with the rest of the urgent data. When finished, TCP informs the application and resumes back to the stream queue. An example is when TCP is used for a remote login session, the user can send a keyboard sequence that interrupts or aborts the program at the other end. These signals are most often needed when a program on the remote machine fails to operate correctly. The signals must be sent without waiting for the program to finish its current transfer.[2]

TCP OOB data was not designed for the modern Internet. The *urgent* pointer only alters the processing on the remote host and doesn't expedite any processing on the network itself. When it gets to the remote host there are two slightly different interpretations of the protocol, which means only single bytes of OOB data are reliable. This is assuming it's reliable at all as it's one of the least commonly used protocol elements and tends to be poorly implemented.[11][12]

[edit] Forcing data delivery

Normally, TCP waits for the buffer to exceed the maximum segment size before sending any data. This creates serious delays when the two sides of the connection are exchanging short messages and need to receive the response before continuing. For example, the login sequence at the beginning of a telnet session begins with the short message "Login", and the session cannot make any progress until these five characters have been transmitted and the response has been received. This process can be seriously delayed by TCP's normal behavior when the message is provided to TCP in several send calls.

However, an application can force delivery of segments to the output stream using a *push* operation provided by TCP to the application layer.[2] This operation also causes TCP to set the PSH flag or control bit to ensure that data is delivered immediately to the application layer by the receiving transport layer.

In the most extreme cases, for example when a user expects each keystroke to be echoed by the receiving application, the *push* operation can be used each time a keystroke occurs. More generally, application programs use this function to force output to be sent after writing a character or line of characters. By forcing the data to be sent immediately, delays and wait time are reduced.

[edit] Connection termination

The connection termination phase uses, at most, a four-way handshake, with each side of the connection terminating independently. When an endpoint wishes to stop its half of the connection, it transmits a FIN packet, which the other end acknowledges with an ACK. Therefore, a typical tear-down requires a pair of FIN and ACK segments from each TCP endpoint.

A connection can be "half-open", in which case one side has terminated its end, but the other has not. The side that has terminated can no longer send any data into the connection, but the other side can. The terminating side should continue reading the data until the other side terminates as well.

It is also possible to terminate the connection by a 3-way handshake, when host A sends a FIN and host B replies with a FIN & ACK (merely combines 2 steps into one) and host A replies with an ACK.[13] This is perhaps the most common method.

It is possible for both hosts to send FINs simultaneously then both just have to ACK. This could possibly be considered a 2-way handshake since the FIN/ACK sequence is done in parallel for both directions.

Some host TCP stacks may implement a "half-duplex" close sequence, as Linux or HP-UX do. If such a host actively closes a connection but still has not read all the incoming data the stack already received from the link, this host sends a RST instead of a FIN (Section 4.2.2.13 in RFC 1122). This allows a TCP application to be sure the remote application has read all the data the former sent—waiting the FIN from the remote side, when it actively closes the connection. However, the remote TCP stack cannot distinguish between a *Connection Aborting RST* and this *Data Loss RST*. Both cause the remote stack to throw away all the data it received, but that the application still didn't read.[*clarification needed*]

Some application protocols may violate the OSI model layers, using the TCP open/close handshaking for the application protocol open/close handshaking - these may find the RST problem on active close. As an example:

```
s = connect(remote);
send(s, data);
close(s);
```

For a usual program flow like above, a TCP/IP stack like that described above does not guarantee that all the data arrives to the other application *unless* the programmer is sure that the remote side will not send anything.

[edit] Vulnerabilities

[edit] Denial of service

By using a spoofed IP address and repeatedly sending purposely assembled SYN packets, attackers can cause the server to consume large amounts of resources keeping track of the bogus connections. This is known as a SYN flood attack. Proposed solutions to this problem include SYN cookies and Cryptographic puzzles. Sockstress is a similar attack, against which no defense is yet known.[citation needed] An advanced DoS attack involving the exploitation of the TCP Persist Timer was analyzed at Phrack #66.[14]

[edit] Connection hijacking

Main article: TCP sequence prediction attack

An attacker who is able to eavesdrop a TCP session and redirect packets can hijack a TCP connection. To do so, the attacker learns the sequence number from the ongoing communication and forges a false segment that looks like the next segment in the stream. Such a simple hijack can result in one packet being erroneously accepted at one end. When the receiving host acknowledges the extra segment to the other side of the connection, synchronization is lost. Hijacking might be combined with ARP or routing attacks that allow taking control of the packet flow, so as to get permanent control of the hijacked TCP connection.[15]

Impersonating a different IP address was possible prior to RFC 1948, when the initial *sequence number* was easily guessable. That allowed an attacker to blindly send a sequence of packets that the receiver would believe to come from a different IP address, without the need to deploy ARP or routing attacks: it is enough to ensure that the legitimate host of the impersonated IP address is down, or bring it to that condition using denial of service attacks. This is why the initial sequence number is chosen at random.

[edit] TCP ports

Main article: TCP and UDP port

TCP uses the notion of port numbers to identify sending and receiving application end-points on a host, or *Internet sockets*. Each side of a TCP connection has an associated 16-bit unsigned port number (0-65535) reserved by the sending or receiving application. Arriving TCP data packets are identified as belonging to a specific TCP connection by its sockets, that is, the combination of source host address, source port, destination host address, and destination port. This means that a server computer can provide several clients with several services simultaneously, as long as a client takes care of initiating any simultaneous connections to one destination port from different source ports.

Port numbers are categorized into three basic categories: well-known, registered, and dynamic/private. The well-known ports are assigned by the Internet Assigned Numbers Authority (IANA) and are typically used by system-level or root processes. Well-known applications running as servers and passively listening for connections typically use these ports. Some examples include: FTP (21), SSH (22), TELNET (23), SMTP (25) and HTTP (80). Registered ports are typically used by end user applications as ephemeral source ports when contacting servers, but they can also identify named services that have been registered by a third party. Dynamic/private ports can also be used by end user applications, but are less commonly so. Dynamic/private ports do not contain any meaning outside of any particular TCP connection.

[edit] Development

TCP is a complex protocol. However, while significant enhancements have been made and proposed over the years, its most basic operation has not changed significantly since its first specification RFC 675 in 1974, and the v4 specification RFC 793, published in September 1981. RFC 1122, Host Requirements for Internet Hosts, clarified a number of TCP protocol implementation requirements. RFC 2581, TCP Congestion Control, one of the most important TCP-related RFCs in recent years, describes updated algorithms that avoid undue congestion. In 2001, RFC 3168 was written to describe explicit congestion notification (ECN), a congestion avoidance signalling mechanism.

The original TCP congestion avoidance algorithm was known as "TCP Tahoe", but many alternative algorithms have since been proposed (including TCP Reno, TCP Vegas, FAST TCP, TCP New Reno, and TCP Hybla).

TCP Interactive (iTCP) [16] is a research effort into TCP extensions that allows applications to subscribe to TCP events and register handler components that can launch applications for various purposes, including application assisted congestion control.

Multipath TCP (MPTCP) [17] is a another research effort attempting to utilize multiple path for one TCP connection, thus maximizing resource usage and increasing redundancy. The redundancy offered by Multipath TCP in the context of wireless networks [18] enables statistical multiplexing of resources, and thus increases TCP throughput dramatically.

TCP Cookie Transactions (TCPCT) is an extension proposed in December 2009 to secure servers against denial-of-service attacks. Unlike SYN cookies, TCPCT does not conflict with other TCP extensions such as window scaling. TCPCT was designed due to necessities of DNSSEC, where servers have to handle large numbers of short-lived TCP connections.

tcpcrypt is an extension proposed in July 2010 to provide transport-level encryption directly in TCP itself. It's designed to work transparently and not require any configuration. Unlike TLS (SSL), tcpcrypt itself does not provide authentication, but provides simple primitives down to the application to do that. As of 2010, the first tcpcrypt IETF draft has been published and implementations exist for several major platforms.

[edit] TCP over wireless networks

TCP has been optimized for wired networks. Any packet loss is considered to be the result of network congestion and the congestion window size is reduced dramatically as a precaution. However, wireless links are known to experience sporadic and usually temporary losses due to fading, shadowing, hand off, and other radio effects, that cannot be considered congestion. After the (erroneous) back-off of the congestion window size, due to wireless packet loss, there can be a congestion avoidance phase with a conservative decrease in window size. This causes the radio link to be underutilized. Extensive research has been done on the subject of how to combat these harmful effects. Suggested solutions can be categorized as end-to-end solutions (which require modifications at the client or server) [19], link layer solutions (such as RLP in CDMA2000), or proxy based solutions (which require some changes in the network without modifying end nodes [19][20].

[edit] Hardware implementations

One way to overcome the processing power requirements of TCP is to build hardware implementations of it, widely known as TCP Offload Engines (TOE). The main problem of TOEs is that they are hard to integrate into computing systems, requiring extensive changes in the operating system of the computer or device. One company to develop such a device was Alacritech.

[edit] Debugging

A packet sniffer, which intercepts TCP traffic on a network link, can be useful in debugging networks, network stacks and applications that use TCP by showing the user what packets are passing through a link. Some networking stacks support the SO_DEBUG socket option, which can be enabled on the socket using setsockopt. That option dumps all the packets, TCP states, and events on that socket, which is helpful in debugging. Netstat is another utility that can be used for debugging.

[edit] Alternatives

For many applications TCP is not appropriate. One big problem (at least with normal implementations) is that the application cannot get at the packets coming after a lost packet until the retransmitted copy of the lost packet is received. This causes problems for real-time applications such as streaming multimedia (such as Internet radio), real-time multiplayer games and voice over IP (VoIP) where it is sometimes more useful to get most of the data in a timely fashion than it is to get all of the data in order.

For both historical and performance reasons, most storage area networks (SANs) prefer to use Fibre Channel protocol (FCP) instead of TCP/IP.

Also for embedded systems, network booting and servers that serve simple requests from huge numbers of clients (e.g. DNS servers) the complexity of TCP can be a problem. Finally, some tricks such as transmitting data between two hosts that are both behind NAT (using STUN or similar systems) are far simpler without a relatively complex protocol like TCP in the way.

Generally, where TCP is unsuitable, the User Datagram Protocol (UDP) is used. This provides the application multiplexing and checksums that TCP does, but does not handle building streams or retransmission, giving the application developer the ability to code them in a way suitable for the situation, or to replace them with other methods like forward error correction or interpolation.

SCTP is another IP protocol that provides reliable stream oriented services similar to TCP. It is newer and considerably more complex than TCP, and has not yet seen widespread deployment. However, it is especially designed to be used in situations where reliability and near-real-time considerations are important.

Venturi Transport Protocol (VTP) is a patented proprietary protocol that is designed to replace TCP transparently to overcome perceived inefficiencies related to wireless data transport.

TCP also has issues in high bandwidth environments. The TCP congestion avoidance algorithm works very well for ad-hoc environments where the data sender is not known in advance, but if the environment is predictable, a timing based protocol such as Asynchronous Transfer Mode (ATM) can avoid TCP's retransmits overhead.

Multipurpose Transaction Protocol (MTP/IP) is patented proprietary software that is designed to adaptively achieve high throughput and transaction performance in a wide variety of network conditions, particularly those where TCP is perceived to be inefficient.

[edit] Checksum computation

[edit] TCP checksum for IPv4

When TCP runs over IPv4, the method used to compute the checksum is defined in RFC 793:

The checksum field is the 16 bit one's complement of the one's complement sum of all 16-bit words in the header and text. If a segment contains an odd number of header and text octets to be checksummed, the last octet is padded on the right with zeros to form a 16-bit word for checksum purposes. The pad is not transmitted as part of the segment. While computing the checksum, the checksum field itself is replaced with zeros.

In other words, after appropriate padding, all 16-bit words are added using one's complement arithmetic. The sum is then bitwise complemented and inserted as the checksum field. A pseudo-header that mimics the IPv4 packet header used in the checksum computation is shown in the table below.

TCP pseudo-header (IPv4)

Bit offset	**0–3**	**4–7**	**8–15**	**16–31**
0	Source address			
32	Destination address			
64	Zeros		Protocol	TCP length
96	Source port			Destination port
128	Sequence number			
160	Acknowledgement number			
192	Data offset	Reserved	Flags	Window
224	Checksum			Urgent pointer
256	Options (optional)			
256/288+	Data			

The source and destination addresses are those of the IPv4 header. The protocol value is 6 for TCP (cf. List of IP protocol numbers). The TCP length field is the length of the TCP header and data.

[edit] TCP checksum for IPv6

When TCP runs over IPv6, the method used to compute the checksum is changed, as per RFC 2460:

> *Any transport or other upper-layer protocol that includes the addresses from the IP header in its checksum computation must be modified for use over IPv6, to include the 128-bit IPv6 addresses instead of 32-bit IPv4 addresses.*

A pseudo-header that mimics the IPv6 header for computation of the checksum is shown below.

TCP pseudo-header (IPv6)

<table>
<tr><th>Bit offset</th><th colspan="2">0 - 7</th><th colspan="2">8–15</th><th colspan="2">16–23</th><th colspan="2">24–31</th></tr>
<tr><td><b>0</b></td><td colspan="8" rowspan="4">Source address</td></tr>
<tr><td><b>32</b></td></tr>
<tr><td><b>64</b></td></tr>
<tr><td><b>96</b></td></tr>
<tr><td><b>128</b></td><td colspan="8" rowspan="4">Destination address</td></tr>
<tr><td><b>160</b></td></tr>
<tr><td><b>192</b></td></tr>
<tr><td><b>224</b></td></tr>
<tr><td><b>256</b></td><td colspan="8">TCP length</td></tr>
<tr><td><b>288</b></td><td colspan="6">Zeros</td><td colspan="2">Next header</td></tr>
<tr><td><b>320</b></td><td colspan="4">Source port</td><td colspan="4">Destination port</td></tr>
<tr><td><b>352</b></td><td colspan="8">Sequence number</td></tr>
<tr><td><b>384</b></td><td colspan="8">Acknowledgement number</td></tr>
<tr><td><b>416</b></td><td>Data offset</td><td>Reserved</td><td colspan="2">Flags</td><td colspan="4">Window</td></tr>
<tr><td><b>448</b></td><td colspan="4">Checksum</td><td colspan="4">Urgent pointer</td></tr>
<tr><td><b>480</b></td><td colspan="8">Options (optional)</td></tr>
<tr><td><b>480/512+</b></td><td colspan="8">Data</td></tr>
</table>

- Source address – the one in the IPv6 header
- Destination address – the final destination; if the IPv6 packet doesn't contain a Routing header, TCP uses the destination address in the IPv6 header, otherwise, at the originating node, it uses the address in the last element of the Routing header, and, at the receiving node, it uses the destination address in the IPv6 header.
- TCP length – the length of the TCP header and data

- Next Header – the protocol value for TCP

[edit] Checksum offload

Many TCP/IP software stack implementations provide options to use hardware assistance to automatically compute the checksum in the network adapter prior to transmission onto the network or upon reception from the network for validation.

[edit] See also

- Connection-oriented protocol
- T/TCP variant of TCP
- TCP and UDP port
- TCP and UDP port numbers for a long list of ports/services
- TCP congestion avoidance algorithms
- Nagle's algorithm
- Karn's Algorithm
- Maximum transmission unit
- IP fragmentation
- Maximum segment size
- Silly window syndrome
- TCP segment
- TCP Sequence Prediction Attack
- SYN flood
- SYN cookies
- TCP Tuning for high performance networks
- Path MTU discovery
- tcphdr - the Unix TCP header structure in the C programming language
- Stream Control Transmission Protocol (SCTP)
- Multipurpose Transaction Protocol (MTP/IP)
- Transport protocol comparison table
- Sockstress

[edit] References

1. ^ Vinton G. Cerf, Robert E. Kahn, *A Protocol for Packet Network Intercommunication*, IEEE Transactions on Communications, Vol. 22, No. 5, May 1974 pp. 637-648
2. ^ a b c d e f g h i j Comer, Douglas E. (2006). *Internetworking with TCP/IP:Principles, Protocols, and Architecture*. **1** (5th ed.). Prentice Hall. ISBN 0130905526.
3. ^ http://www.iana.org/assignments/tcp-parameters/
4. ^ RFC 1323, TCP Extensions for High Performance, Section 2.2
5. ^ RFC 2018, TCP Selective Acknowledgement Options, Section 2
6. ^ RFC 2018, TCP Selective Acknowledgement Options, Section 3
7. ^ RFC 1323, TCP Extensions for High Performance, Section 3.2
8. ^ RFC 1146, TCP Alternate Checksum Options
9. ^ Stone; Partridge (2000). "When The CRC and TCP Checksum Disagree". *Sigcomm*. http://citeseer.ist.psu.edu/stone00when.html
10. ^ http://lwn.net/Articles/92727/

11. ^ Gont, Fernando (2008-11). "On the implementation of TCP urgent data". 73rd IETF meeting. http://www.gont.com.ar/talks/IETF73/ietf73-tcpm-urgent-data.ppt. Retrieved 2009-01-04.
12. ^ Peterson, Larry (2003). *Computer Networks*. Morgan Kaufmann. pp. 401. ISBN 155860832X.
13. ^ Tanenbaum, Andrew S. (2003-03-17). *Computer Networks* (Fourth ed.). Prentice Hall. ISBN 0-13-066102-3.
14. ^ Exploiting TCP and the Persist Timer Infiniteness
15. ^ Laurent Joncheray, *Simple Active Attack Against TCP*, 1995
16. ^ TCP Interactive (iTCP)
17. ^ draft-ietf-mptcp-architecture
18. ^ http://portal.acm.org/citation.cfm?id=1794199
19. ^ *a* *b* "TCP performance over CDMA2000 RLP". http://academic.research.microsoft.com/Paper/3352358.aspx. Retrieved 2010-08-30
20. ^ Muhammad Adeel &

TCP/IP - Transmission Control Protocol / Internet Protocol

TCP/IP is arguably the single most important computer networking technology. The Internet and most home networks support TCP/IP as the "language" computers use to find and connect with each other.

7. Working with IP Addresses (46)
   8. Working With MAC Addresses (7)
   9. FTP File Transfer (8)
7. TCP/IP Benchmark Tools (5)
8. IP and IPv6 Protocols @
9. DNS - Domain Name System @

TCP/UDP Port Numbers

IP networking uses the concept of protocol "ports" at the transport layer to manage communication channels between programs. Many TCP/IP port numbers are associated with specific applications as explained here.

What Is an IP Address?
An IP address is a logical address of a network adapter. Home computers and many other devices acquire IP addresses to join Internet Protocol networks

Configure or Change IP Address Information

You need to have IP addresses set up correctly when configuring your network, and you may also need to change them occasionally.

How To Change Your IP Address

With some effort, its possible to change your IP address. The procedure depends on whether the IP address is static or dynamic and also whether it is public or private

Release and Renew IP Addresses on Windows XP | 2000 | NT

Follow these step-by-step instructions to quickly release and/or renew Internet Protocol (IP) addresses of Windows XP, Windows 2000, or Windows NT computers.

How to disable network bindings using the [Netbindings] section

Windows 2000 unattended Setup has a new section in the answer file that allows you to disable network bindings on the network card during an unattended Setup. However, after using this section in the unattended Setup file, the bindings appear to be unaffected and still enabled.

This is due to incorrect parameters specified in the [Netbindings] section. The Unattend.doc has incorrect information on the format of the [Netbindings] entries. There should not be any commas (**,**) between the entries, and not all bindings paths follow the example listed in the Unattend.doc.

The Ties That Bind

Network Bindings Help Us All Just Get Along

"Binding as a noun is something on your skis," says Harry Brelsford, author of the MCSE (Microsoft Certified Systems Engineer) Consulting Bible and practicing MCSE consultant. "In computer networking, binding is a verb. It's like glue. You bind a protocol suite to a network interface."

Most people would rather worry about skis than tackle the delicacies of effective networking. However, if you're responsible for a network of any size or are a network user in need of help maximizing your machine's performance, you should grasp the fundamentals of binding protocols and how to manipulate them.

Bindings Explained

Put simply, networks can take on a dizzying number of topologies and configurations. Some LANs (local-area networks) work with Microsoft Windows and networks built for Novell NetWare. Some use Ethernet protocols, while others use Token Ring technology. You can network the computers within a single building or multiple buildings, or you can connect your PC to the world's largest computer network, the Internet. Your computer may be a server or a client. Your NIC (network interface card) may be an Ethernet adapter, a dial-up modem, or a virtual adapter, as with VPNs (Virtual Private Networks). Whatever the case, one computer should be able to network in all these conditions.

Interoperability is key. One system should be able to support and run any number of network protocols. Simultaneously, that computer should be able to run a bevy of network adapters, the devices that physically link the system with others in the network, plus run multiple services for that network, such as file and printer sharing. The trick is to tell each adapter not only to what it will connect, but also which protocol it will use when making that connection. This association of adapter to protocol and protocol to service is the binding process.

Why do you need to understand bindings? Here's a real-life example: A home business owner subscribes to a broadband satellite Internet service, binds his modem/network adapter to the TCP/IP (Transmission Control Protocol/Internet Protocol) protocols, and everything works just fine. Then one day he decides to network his three home PCs and binds his three Ethernet cards to the NetBEUI (Network Basic Input/Output System Extended User Interface, pronounced "NET-boo-ee") protocol. Suddenly, his Internet connection dies. Diagnostic software reports that the modem's USB (Universal Serial Bus) connection has become unplugged, but no amount of replugging or rebooting fixes the problem. The real problem turns out to be that when he bound his Ethernet cards to NetBEUI, Windows automatically bound NetBEUI back to his satellite modem, resulting in a protocol conflict that killed his connection. With a basic grasp of network bindings, the fledgling systems administrator knew to pull up the properties for his modem and unbind NetBEUI. This left his in-house network intact and immediately re-established his Internet connection.

The Protocols

Before moving on, you should understand the role of protocols in networking. A network protocol is like a set of rules for communicating. When you meet your company's president in the boardroom, you shake hands, use his or her title, and say something like "Good afternoon." When you meet a co-worker at the company picnic, you're still likely to shake hands, but you'll probably drop titles, perhaps switching to nicknames, and say something like, "How's it going?" You might call these greeting protocols, wherein you use a different one for each social network.

Computer networks use similar rule sets to pass information between components and machines, and different networks may require different protocols. This may lead you to wonder why network engineers don't just take a USB-style approach and let Windows automatically support all possible protocols, and bind them as needed.

"Automatic detection would introduce a huge amount of overhead on your system," says Ward Ralston, senior staff instructor at Paladin Data Systems. "There are literally thousands of protocols out there, and to have your system automatically detect different kinds of protocols would make things very slow."

Fortunately, although there *are* thousands of network protocols out there, there are only a few you'll need to worry about where networked desktop PCs are concerned. To help you better understand exactly what you're binding to your network adapters, we'll describe these protocols here.

The Advanced Settings dialog box lets you use up- and down-arrow buttons to rearrange the order of your bindings for more efficient

communication.

TCP/IP. This is the protocol your computer uses to communicate with others on the Internet. Originally designed by the U.S. government during the Internet's infancy, the vendor-neutral protocol offers numerous features for ensuring the integrity of the **packets**, or pieces of data, traveling back and forth. It also makes certain that the network adapter and computer on the receiving end correctly reassemble them into whole files. Because TCP/IP is a routable technology, it's ideal for large networks, such as the Internet and corporate WANs (wide-area networks). TCP/IP is less commonly used for smaller, in-house networks, although acceptance is slowly growing. Because of its complexity, TCP/IP isn't among the fastest or smallest protocols in terms of resource overhead, but modern hardware and software have done much to negate the protocol's drawbacks.

IPX/SPX (Internetwork Packet Exchange/Sequenced Packet Exchange). Descended from the XNS (Xerox Network Systems) protocol, Novell optimized IPX for Ethernet networks, although it also works with alternative topologies such as Token Ring. Novell designed it specifically for use with LANs, a field the company dominated throughout most of the ‘80s and ‘90s. Recent versions of Windows support IPX/SPX with help from NWLink, a Microsoft protocol that negates the need for Novell's file and printer sharing software. IPX/SPX protocols require recipient computers to acknowledge receipt of data packets, which helps ensure data integrity, but can drag down performance due to increased network traffic. Recent enhancements to the specification let it use larger packet sizes to decrease the total number of packets sent.

NetBEUI. NetBEUI is the closest thing around to a native network protocol for Windows-based PCs. NetBIOS began as a set of 18 commands for networking IBM PCs, and NetBEUI built on this foundation. NetBEUI is not a routable protocol, making it unsuited to WANs and other large networks, but it's very fast, has strong error correction, and is simple to set up. The protocol is compatible with all Windows versions starting with Windows for Workgroups, and is often the protocol of choice for small Windows networks.

All systems on a network must use the same protocol, but that doesn't mean that each can only run one protocol at a time. A server might connect to a workgroup across the hall using IPX/SPX bound to a 3Com card, an office on the next floor via NetBEUI on a NETGEAR adapter, and to the ‘Net with TCP/IP bound to yet another card—if you manage the bindings properly.

The Network Bindings Dialog Box(es)

The dialog boxes you'll use to work with bindings vary from OS (operating system) to OS. In general, though, all binding controls in Windows are available through the Network area of Control Panel. These boxes let you see how your services, protocols, and adapters are bound together; let you bind and unbind components to one another; and establish the order in which your PC addresses protocols.

Occasionally, you might hear that binding order is irrelevant to network operation, based on the fact that when you add a new protocol, Windows may relegate your active binding to the bottom of the binding list. However, binding order can be important. When sending data, the computer attempts communication across each binding, continuing the search until it finds the correct communication channel. Such delays won't make much difference to home users but can have a profound effect on large, corporate networks.

To change your binding order within Windows 2000 Professional, click Start, point to Settings, click Control Panel, and double-click Network And Dial-Up Connections. (You can also right-click My Network Places on your Desktop and select Properties.) Click your network adapter, then click File, point to Advanced, and click Advanced Settings. This brings up the Advanced Settings dialog box, and you should be looking at the Adapters And Bindings tab. Under Client For Microsoft Networks, there are three bound protocols: NWLink, NetBEUI, and TCP/IP. This adapter connects a PC to its broadband Internet connection, so you should unbind NWLink, then highlight Internet Protocol (TCP/IP). Using the up-arrow button near the Bindings box, click twice to move TCP/IP to the top of the list, then click OK to save the new settings.

Another way to unbind protocols is to open Network And Dial-Up Connections, right-click your adapter, and select Properties. You'll see a list of all the bound services and protocols associated with that adapter in the Properties dialog box. Deselect all but the essential components, then click Close to save the changes. You can click the Uninstall button to remove a component's binding to your adapter, but doing so will unbind the component from all adapters, so it's best to unbind on a case-by-case basis.

■Step Through The Wizard

Let's say you want to create a network connection for moving files, such as a peer-to-peer parallel connection. In most versions of Windows, including Windows Me, you would right-click Network Neighborhood/My Network Places, select Properties, and see a dialog box offering your choice of components to install. Windows 2000 Professional offers a more spoon-fed approach with its Network Connection Wizard. Click Start, Settings, Network And Dial-Up Connections, and Make New Connection. Click Next at the Welcome screen, Select Accept Incoming Connections, and click Next.

Windows should detect (at least) two suitable adapters for this operation on most PCs: a dial-up modem and a parallel port. Select the latter and click Next. For this example you won't need the security or extra complexity of a VPN, so opt not to allow for an incoming VPN connection and proceed. Select users you want to give access to your connection from the list that appears, and click Next.

In the Networking Components dialog box, you'll find a list of network services and protocols from which to choose. If the component you want isn't on the list, click the Install button (you may be prompted to insert a suitable disc for the source files). Be sure to unbind any components you don't want to apply to this connection. Click Next, pick a name for your new connection, click Finish, and you're done.

■Final Word

Once you grasp the principles of binding services to protocols and protocols to adapters, the binding process becomes simple and useful. Armed with this knowledge, you can be sure of setting up and maintaining a well-bound network that will last you and your co-workers for years. ■

by William Van Winkle

Binding Up The Security Leak

One of the most infamous security holes in Windows-based PCs is the File And Printer Sharing service. Ward Ralston, senior staff instructor at Paladin Data Systems, has his class use a free sniffer program (one that monitors data traffic to and from various points on a network) to examine a block of cable modem IP (Internet Protocol) addresses as a security management exercise. Out of about 16,000 addresses, he says, generally 500 or 600 have the service enabled and are bound to the TCP/IP (Transmission Control Protocol/Internet Protocol) protocol. This gives anyone full access to the users' systems over the Internet, unless they have set up sharing restrictions and permisions.

Microsoft enables File And Printer Sharing automatically under several circumstances. Poorly configured broadband connections are one common example, but most home networking kits will also turn on the feature. Say you install a kit that binds NWLink to your network card. In order to transfer files between PCs, the kit automatically installs the File And Printer Sharing service, then binds it with the NWLink protocol to assist communication across the small network. The problem is that Windows defaults to binding all possible services and protocols to a new component when you add one. So, in addition to binding NWLink to File And Printer Sharing, Windows binds it to the already-installed TCP/IP, which it might already have bound to the dial-up adapter. Suddenly, hackers have an open front door to your system.

Some sources simply advocate disabling File And Printer Sharing, but this would negate having a LAN. Here's one answer: Back in your Network Bindings options area (Advanced Settings in Windows 2000), look at File And Printer Sharing For Microsoft Networks. Under this you should see at least two bound protocols: NWLink IPX/SPX/NetBIOS Compatible Transport Protocol and Internet Protocol (TCP/IP). Unbind TCP/IP. This lets NWLink remain active (so your LAN stays operational), but prohibits any TCP/IP-based Internet traffic from getting past your network adapter. This will cut your client PCs off from the Internet, but they'll still have full file sharing capabilities. In fact, you might use the isolated PCs to store sensitive information until you can implement a more permanent solution, such as a firewall.

How to configure TCP/IP to use DNS in Windows XP

This article describes how to configure Windows XP TCP/IP to use Domain Name Service (DNS).

Back to the top

How to configure TCP/IP

To configure TCP/IP, follow these steps:

1. Click **Start**, click **Control Panel**, click **Network and Internet Connections**, and then click **Network Connections**.
2. Right-click the network connection that you want to configure, and then click **Properties**.
3. On the **General** tab (for a local area connection), or the **Networking** tab (for all other connections), click **Internet Protocol (TCP/IP)**, and then click **Properties**.

4. If you want to obtain DNS server addresses from a DHCP server, click **Obtain DNS server address automatically**.
5. If you want to manually configure DNS server addresses, click **Use the following DNS server addresses**, and then type the preferred DNS server and alternate DNS server IP addresses in the **Preferred DNS server** and **Alternate DNS server** boxes.

Back to the top

How to configure the advanced DNS properties

To configure the advanced DNS properties, follow these steps:

1. Click **Advanced**, and then click the **DNS** tab.
2. Configure the DNS properties by following the steps that are described in the following sections:
    - How to configure an additional DNS server IP address
    - How to modify the resolution behavior for unqualified DNS names
    - How to modify DNS dynamic update behavior

How to configure an additional DNS server IP address

To configure an additional DNS server IP address, follow these steps:

1. Under **DNS server addresses**, click **Add** in order of use.
2. In the **TCP/IP DNS server** box, type the IP address of the DNS server, and then click **Add**.

How to modify the resolution behavior for unqualified DNS names

To modify the resolution behavior for unqualified DNS names, follow these steps:

1. To resolve an unqualified name by appending the primary DNS suffix and the DNS suffix of each connection, click **Append primary and connection specific DNS suffixes**. To do this, each connection must be configured. If you also want to search the parent suffixes of the primary DNS suffix up to the second-level domain, click to select the **Append parent suffixes of the primary DNS suffix** check box.
2. To resolve an unqualified name by appending the suffixes from a list of configured suffixes, click **Append these DNS suffixes (in order)**, and then click **Add** to add suffixes to the list.
3. To configure a connection-specific DNS suffix, type the DNS suffix in the **DNS suffix for this connection** box.

How to modify DNS dynamic update behavior

To modify DNS dynamic update behavior, use any of the following methods:

- To use a DNS dynamic update to register the IP addresses of this connection and the primary domain name of the computer, select the **Register this connection's addresses in DNS** check box. By default, this check box is not selected. The primary domain name of the computer is the primary DNS suffix appended to the computer name. To view this domain name and DNS suffix together as the full computer name, click **Start**, click **Control Panel**, click **Performance and Maintenance**, click **System**, and then click the **Computer Name** tab.
- To use a DNS dynamic update to register the IP addresses and the connection-specific domain name of this connection, select the **Use this connection's DNS suffix in DNS registration** check box. By default, this check box is not selected. The connection-specific domain name of this connection is the DNS suffix for this connection appended to the computer name.

↑Back to the top

Troubleshooting

To disable DNS dynamic update for all names on the computer, clear the **Register this connection's addresses in DNS** and the **Use this connection's DNS suffix in DNS registration** check boxes on the **DNS** tab for all connections in **Network Connections**.

Note You must be logged on as an administrator or as a member of the Administrators group to use this procedure. If your computer is connected to a network, network policy settings may also prevent you from using this procedure.

Internet Protocol security (IPsec) is a framework of open standards for protecting communications over Internet Protocol (IP) networks through the use of cryptographic security services. IPsec supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection. The Microsoft implementation of IPsec is based on standards developed by the Internet Engineering Task Force (IETF) IPsec working group.

IPsec is supported by the Microsoft Windows 7, Windows Server 2008 R2, Windows Vista, Windows Server 2008, Windows Server 2003, Windows XP, and Windows 2000 operating systems and is integrated with Active Directory Domain Services (AD DS). IPsec policies can be assigned through Group Policy, which allows IPsec settings to be configured at the domain, site, organizational unit, or security group level.

In Windows 7, Windows Server 2008 R2, Windows Vista and Windows Server 2008, you can configure IPsec behavior with the Windows Firewall with Advanced Security snap-in.

IPSec troubleshooting tools

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Troubleshooting tools

This topic provides information about the following IPSec troubleshooting tasks and the troubleshooting tools that you can use to perform each of these tasks:

- Viewing IPSec policy assignment information
- Viewing details about the active IPSec policy and IPSec statistics
- Viewing details about the active IPSec policy and IPSec statistics in IP Security Monitor
- Verifying that security auditing is enabled
- Viewing IPSec-related events in Event Viewer
- Enabling detailed tracing for Internet Key Exchange (IKE) negotiations
- Viewing IPSec and other network communication with Network Monitor
- Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family
- Using Ipseccmd.exe to manage and monitor IPSec on computers running Windows XP
- Using Netdiag.exe to display IPSec information and to test and view network configuration
- Disabling TCP/IP and IPSec hardware acceleration

Details about each of these subjects are provided in the following sections.

Notes

- For general recommended practices that can help you enhance security and minimize the potential for problems when deploying IPSec, see IPSec Best practices.
- For the Windows Server 2003 family, the Netsh commands for IPSec replace and enhance the functionality provided by Netdiag.exe and Ipseccmd.exe. For information about how to use Netsh to script IPSec policy creation and to monitor IPSec activity, see Netsh commands for Internet Protocol security.
- Ipseccmd.exe is a support tool that is provided only on the Windows XP CD. To install this tool, run the Setup.exe file from the Support\Tools folder. You can use this tool to manage and monitor IPSec policies only on computers running Windows XP. For more information about Ipseccmd.exe, see the Microsoft Tech Net Web site.

Netdiag.exe is available in the Windows Server 2003 family, Windows XP, and in Windows 2000, as follows:

- In the Windows Server 2003 family, run the Suptools.msi file from the \Support\Tools folder on the Windows Server 2003 family CD (select the **Complete** setup option). Although you can still use this version of Netdiag.exe to obtain basic networking information that is not IPSec-specific, the Netsh commands for IPSec replace all IPSec-specific functionality.
- In Windows XP, run the Setup.exe file from the \Support\Tools folder on the Windows XP CD (select the **Complete** setup option).

- For Windows 2000, an updated version of Netdiag.exe is available for download from the Web. For more information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site.

For more information about Netdiag.exe, see Using Netdiag.exe to display IPSec information and to test and view network configuration in this topic.

Viewing IPSec policy assignment information

Viewing the name of an active IPSec policy and the name of the Group Policy object to which the active IPSec policy is assigned can be useful for troubleshooting policy precedence issues. The following table summarizes the tools that you can use to view this information for the Windows Server 2003 family, Windows XP, and Windows 2000.

On computers running	Use these tools to view the name of the active IPSec policy	Use these tools to view the name of the Group Policy object to which the active IPSec policy is assigned
Windows Server 2003 family	• IP Security Monitor console. For more information, see View details about active IPSec policies. • Netsh, **netsh ipsec static show gpoassignedpolicy** command. For more information, see Netsh commands for Internet Protocol security.	• Resultant Set of Policy (RSoP) console. For more information, see Use Resultant Set of Policy (RSoP) to View IPSec Policy Assignments. • Netsh, **netsh ipsec static show gpoassignedpolicy** command.
Windows XP	• IP Security Policy Management console (for local policies only). • Network Connectivity Tester (Netdiag.exe), **netdiag /test:ipsec** command. For more information, see Using Netdiag.exe to display IPSec information and to test and view network configuration in this topic.	• Netdiag.exe, **netdiag /test:ipsec** command.
Windows 2000	• Netdiag.exe, **netdiag /test:ipsec** command. • In the properties for the relevant network connection, **TCP/IP**	• Netdiag.exe, **netdiag /test:ipsec** command. • Gpresult.exe (Group

	Properties/Advanced/Options/IPSec. The assigned IPSec policy that is displayed in **TCP/IP Properties** is global. It is not specific to the connection.	Policy Results). • Gpotool.exe (Group Policy Verification Tool). Gpresult.exe and Gpotool.exe are both available for download from the Web. For more information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site.

Notes

- To view all IPSec policies that are available (but not necessarily assigned or applied) to computers, use the IP Security Policy Management console. For information about policy precedence and IPSec policy behavior in an Active Directory environment, see the section "Active Directory-based policy" in Creating, modifying, and assigning IPSec policies.
- For the Windows Server 2003 family, to determine which IPSec policies are assigned but are not being applied to IPSec clients, use the RSoP console.
- The Windows XP implementation of the RSoP console does not support the display of IPSec policies. In addition, the **gpresult /scope computer** command does not display the Group Policy object that contains an IPSec policy assignment. For these reasons, you should use Netdiag.exe to view IPSec policy assignment information on computers running Windows 2000 or Windows XP. The **netdiag /test:ipsec** command displays the Group Policy object that contains the IPSec policy assignment, and the organizational unit to which the Group Policy object is assigned.
- Gpotool.exe allows you to monitor the health of Group Policy objects on domain controllers. You can use this tool to check the consistency and replication of Group Policy objects and to display Group Policy object properties. Gpotool.exe is available only for computers running Windows 2000.

Viewing details about the active IPSec policy and IPSec statistics

The IP Security Policy Management console, Netdiag.exe, Gpresult.exe, and Gpotool.exe allow you to determine which IPSec policy has been assigned through Group Policy. After you verify this information, you might need to view details about the assigned IPSec policy and IPSec statistics (for example, filters, filter actions, and active security associations). The following table summarizes the tools that you can use to view an active IPSec policy and IPSec statistics for the Windows Server 2003 family, Windows XP, and Windows 2000.

On computers running	Use these tools to display the active IPSec policy, IPSec statistics, or both	Notes
Windows Server 2003 family	• IP Security Monitor console.	• To view details about the active IPSec policy and IPSec statistics on a local or remote computer, you must be a member of the

	For more information, see View IP security statistics • Netsh, **netsh ipsec dynamic show all** command.	Administrators group on that computer.
Windows XP	• IP Security Monitor console. • IPseccmd.exe, **ipseccmd show all** command.	• To view details about the active IPSec policy and IPSec statistics for a local or remote computer, you must be a member of the Administrators group on that computer.
Windows 2000	• Netdiag.exe, **netdiag /test:ipsec /v /debug** command. • Ipsecmon.exe.	• To use the **/debug** option, you must be logged on as a member of the Administrators group on that computer. In addition, to view details about Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory. • Ipsecmon.exe only displays active outbound quick mode security associations (SAs).

Note

- To remotely monitor IPSec on a computer that is running a different version of Windows than your computer, use Remote Desktop Connection. For example, if your computer is running the Windows Server 2003 family and you plan to remotely monitor IPSec on computers running Windows 2000 or Windows XP, use Remote Desktop Connection to gain remote access to these computers. For information about Remote Desktop Connection, see Remote Desktop Connection.

Viewing details about the active IPSec policy and IPSec statistics in IP Security Monitor

In Windows XP and the Windows Server 2003 family, IP Security Monitor is implemented as a Microsoft Management Console (MMC) snap-in, and it includes enhancements that allow you to view details about an active IPSec policy that is applied by the domain or locally, as well as quick mode and main mode statistics, and active IPSec SAs. IP Security Monitor also enables you to search for specific main mode or quick mode filters. To troubleshoot complex IPSec policy designs, you can use IP Security Monitor to search for all matches for filters of a specific traffic type.

- For procedures on how to use IP Security Monitor, see Monitor IPSec Activity.
- For detailed information about IP security statistics, see Viewing main mode and quick mode statistics in IP Security Monitor.

Notes

- You can use IP Security Monitor to monitor only computers that are running Windows XP or the Windows Server 2003 family. In addition, the computer that is being monitored must run the same version of the Windows operating system as the computer on which IP Security Monitor is running. To monitor IPSec on a computer that is running Windows 2000, use the **ipsecmon** command at the Windows 2000 command prompt on the computer that is being monitored.

- For remote monitoring, you can use IP Security Monitor only to monitor computers that are running the same version of the Windows operating system. To remotely monitor IPSec on a computer that is running a different version of Windows than your computer, use Remote Desktop Connection. For information about Remote Desktop Connection, see Remote Desktop Connection.

- If your computer is running the Windows Server 2003 family, and you plan to monitor IPSec on computers that are also running the Windows Server 2003 family, you can run IP Security Monitor or use the Netsh command-line tool remotely.

Verifying that security auditing is enabled

You can use Local Security Policy settings (for a local computer) or Group Policy Object Editor (for a domain) to verify that security auditing is enabled, so you can ensure that the success and failure of IKE negotiations is recorded. Auditing for IKE is supported in Windows 2000, Windows XP, and the Windows Server 2003 family (IKE uses the Logon Events category). In the Windows Server 2003 family, you can also enable auditing for the security policy database (SPD). SPD uses the Policy Change category. For more information, see Define or modify auditing policy settings for an event category.

Viewing IPSec-related events in Event Viewer

You can use Event Viewer to view the following IPSec-related events:

IKE events (negotiation success and failure) in the security log.

To view these events, enable success or failure auditing for the **Audit logon events** audit policy for your domain or local computer. The IKE event category is also used for auditing user logon events in services other than IPSec. For more information, see Define or modify auditing policy settings for an event category.

When you enable success or failure auditing for the **Audit logon events** audit policy, IPSec records the success or failure of each main mode and quick mode negotiation and the establishment and termination of each negotiation as separate events. However, enabling this type of auditing can cause the security log to fill with IKE events. For example, for servers that are connected to the Internet, attacks on the IKE protocol can cause the security log to fill with IKE events. IKE events can also fill the security log for servers that use IPSec to secure traffic to many clients. To avoid this, you can disable auditing for IKE events in the security log by modifying the registry.

- To disable auditing of IKE events in the security log, do the following:

  1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Audit\DisableIKEAudits** registry setting to a value of **1**.

     The **DisableIKEAudits** key does not exist by default and must be created.

Caution

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

3. Restart the computer, or stop and then restart the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

Note

- Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. If you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established. For more information, see "Stopping and restarting the IPSec service," later in this topic.

IPSec policy change events in the Security log.

To view these events, enable success or failure auditing for the **Audit policy change** audit policy for your domain or local computer. For more information, see Define or modify auditing policy settings for an event category.

IPSec driver per-packet drop events in the System log.

In Windows 2000, Windows XP, and the Windows Server 2003 family, you can enable packet event logging for the IPSec driver by modifying the registry. The IPSec driver reads the registry during computer startup. In the Windows Server 2003 family, you can enable packet event logging for the IPSec driver by using the **Netsh ipsec** command-line tool. To enable logging of dropped inbound and outbound packets, specify a value of **7**. For information about the other levels of IPSec driver event logging, see Notes.

- To enable IPSec driver logging of dropped inbound and outbound packets by modifying the registry, do the following:

    1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec\EnableDiagnostics DWORD** registry setting to a value of **7**.

Caution

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

3. Restart the computer.

- To enable IPSec driver logging of dropped inbound and outbound packets by using the Netsh IPSec command-line tool in the Windows Server 2003 family, do the following:

    0. At the command prompt, type: **netsh ipsec dynamic set config ipsecdiagnostics 7**

    1. Restart the computer.

You can also change the interval for writing IPSec driver packet events to the System log. By default, the IPSec driver writes events to the System log once an hour or after a threshold for the number of events has been reached. For troubleshooting, you should set this interval to the minimum value, 60 seconds.

- To change the interval for writing IPSec driver packet events by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec\LogInterval DWORD** registry setting to **60 decimal**.

Caution

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

3. Restart the computer.

- To change the interval for writing IPSec driver packet events by using the **Netsh ipsec** command-line tool in the Windows Server 2003 family, do the following:

0. At the command prompt, type **netsh ipsec dynamic set config ipsecloginterval 60**

1. Restart the computer.

The IPSec driver reads the registry during computer startup.

By default, packet event logging for the IPSec driver is disabled (that is, the value is set to **0**). When you enable packet event logging for the IPSec driver, you can specify any of the following values, to enable different levels of logging:

Logging Level	Description
1	When **1** is specified, bad SPI packets (the total number of packets for which the Security Parameters Index or SPI was incorrect), IKE negotiation failures, IPSec processing failures, packets received with invalid packet syntax, and other errors are recorded in the System log. Unauthenticated hashes (with the exception of the "Clear text received when should have been secured" event) are logged as well.
2	When **2** is specified, inbound per-packet drop events are recorded in the System log.
3	When **3** is specified, level 1 and level 2 logging are performed. In addition, unexpected clear text events (packets that are sent or received in plaintext) are also recorded.
4	When **4** is specified, outbound per-packet drop events are recorded in the System log.
5	When **5** is specified, level 1 and level 4 logging are performed.
6	When **6**is specified, level 2 and level 4 logging are performed.
7	When **7** is specified, all levels of logging are performed.

For more information about IPSec driver event logging, see the Microsoft Windows Resource Kits Web site.

You cannot configure audit policies on a computer running Windows XP Home Edition. However, success and failure auditing for the **Audit logon events** and **Audit policy change** audit policies for the local computer are enabled by default.

Enabling auditing in the security log or IPSec driver diagnostics in the system log can cause these logs to fill with events quickly. Before you perform either of these tasks, you should do the following:

- Ensure that the size of the log is at least 10 megabytes (MB).
- Save the existing log to a file.
- Clear all events in the log so that the log is empty.
- Consider disabling auditing of IKE events in the security log by modifying the registry, if your computer hosts many simultaneous network connections.

Enabling detailed tracing for Internet Key Exchange (IKE) negotiations

Enabling audit logging for IKE events and viewing the events in Event Viewer is the fastest and simplest way to troubleshoot failed main mode or quick mode negotiations. However, some scenarios might require a more detailed analysis of the IKE main mode negotiation and quick mode negotiations for troubleshooting. You can enable tracing for IKE negotiations if the audit failure events do not provide enough information. The IKE tracing log is a very detailed log intended for troubleshooting IKE interoperability under controlled circumstances. Expert knowledge of the ISAKMP RFC 2408 and IKE RFC 2409 is required to interpret this log.

The IKE tracing log appears as the *systemroot*\Debug\Oakley.log file. The log has a fixed size of 50,000 lines and will overwrite as necessary. A new Oakley.log file is created each time the IPSec service is started and the previous version of the Oakley.log file is saved as Oakley.log.sav. When the Oakley.log file becomes full, the current file is saved as Oakley.log.bak, and a new Oakley.log file is created.

Because many IKE negotiations can occur simultaneously, you should minimize the number of negotiations and log for as short a period of time as possible to capture a more easily interpreted log.

Enabling and disabling the IKE tracing log in the Windows Server 2003 **family**

In the Windows Server 2003 family, you can enable or disable the IKE tracing log dynamically while the IPSec service is running by doing the following:

- To enable the IKE tracing log, type the following at the command prompt:

  netsh ipsec dynamic set config ikelogging 1

  This command creates the IKE tracing log file if it does not exist. If the file does exist, it appends logging information to the existing file.

- To disable the IKE tracing log, type the following at the command prompt:

  netsh ipsec dynamic set config ikelogging 0

Enabling the IKE tracing log in Windows 2000 and Windows XP

In Windows 2000 and Windows XP, you must enable IKE tracing by modifying the registry. For the changes to take effect, you must also stop and restart the IPSec service:

- To enable the IKE tracing log in Windows XP and Windows 2000, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLoggingDWORD** registry setting to a value of **1**.

   The Oakley key does not exist by default and must be created.

Caution

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

3. Stop and start the IPSec service by running the **net stop policyagent** and **net start policyagent** commands at the command prompt.

Stopping and restarting the IPSec service

Stopping and restarting the IPSec service can disconnect all of the computers that are using IPSec from the computer on which the IPSec service is stopped, and it can prevent further communication with that computer. However, if you restart the IPSec service immediately, TCP-based communication might resume, due to the retransmit behavior of TCP, after new IKE and IPSec SAs are established. If you leave the IPSec service stopped, any clients that used the default response rule to establish security will be unable to communicate with this computer for two hours.

To avoid losing Terminal Services connectivity for computers that are using IPSec over a Terminal Server session, you must stop and restart the IPSec service by using a single command line:

- To stop and restart the IPSec service for computers over a Terminal Server session

  - At the command prompt, type the following:

    net stop policyagent & net start policyagent

If you are restarting the IPSec service while an L2TP/IPSec VPN tunnel is connected, the tunnel will lose connectivity and must be reconnected when the IPSec service is restarted. A PPTP tunnel, however, does not use IPSec and therefore will stay connected if you stop the IPSec service.

If you are restarting the IPSec service on a computer that is running the Windows Server 2003 family or the Windows Server 2003 family and that is also running the Routing and Remote Access service, any IPSec configuration for L2TP will be lost and the L2TP tunnels will be disconnected. Therefore, you must stop and restart the Routing and Remote Access service, as well as the IPSec service.

- To stop and restart the IPSec service and the Routing and Remote Access service, do the following:

  1. Stop the Routing and Remote Access service using the **net stop remoteaccess** command.

  2. Stop the IPSec service by using the **net stop policyagent** command.

  3. Start the IPSec service by using the **net start policyagent** command.

  4. Start the Routing and Remote Access service using the **net start remoteaccess** command.

Viewing IPSec and other network communication with Network Monitor

You can install and use Network Monitor to view IPSec and other network communication. Note that the version of Network Monitor that is provided with the Windows Server 2003 family can be used only to view the network traffic that is sent to or from the computer on which it is installed. To view network traffic that is sent to or from another computer and is routed through your computer (using the Routing and Remote Access service), you must use the Network Monitor component that is provided with Microsoft Systems Management Server.

The Network Monitor component that is provided with the Windows Server 2003 family includes parsers for the ISAKMP (IKE), AH, and ESP protocols. The Network Monitor parsers for ESP can parse inside the ESP packet only if null-encryption is being used and the full ESP packet is captured. Network Monitor cannot parse the encrypted portions of IPSec-secured ESP traffic when encryption is performed in software. However, if encryption is being performed by an IPSec hardware offload network adapter, the ESP packets are decrypted when Network Monitor captures them and as a result, can be parsed and interpreted into the upper-layer protocols. If you need to diagnose ESP software-encrypted communication, you must disable ESP encryption and use ESP-null encryption by changing the IPSec policy on both computers.

Notes

- For information about how to install the version of Network Monitor that is provided with the Windows Server 2003 family, see Install Network Monitor.
- For information about Systems Management Server, go to the Microsoft Web site.

Using Netsh to change the IPSec configuration on computers running the Windows Server 2003 family

For computers running the Windows Server 2003 family, you can use the Netsh commands for IPSec to script IPSec policy creation, display details about IPSec policies, and change the IPSec configuration for troubleshooting.

For troubleshooting, you can use the Netsh commands for IPSec to change the following settings:

Disable Internet Key Exchange (IKE) certificate revocation list (CRL) checking
By default, in Windows 2000 CRLs are not checked during IKE certificate authentication. In Windows XP and the Windows Server 2003 family, CRLs are checked during IKE certificate authentication, but a fully successful check is not required for the certificate to be accepted. In some cases, failures during CRL processing might cause IKE to not accept the certificate. Or, the delay required for CRL checking might delay IKE negotiation enough to cause the connection attempt to time-out. To determine whether certificate authentication will be successful without CRL checking, you can disable IKE CRL checking. To do this, type the following at the command prompt:

```
netsh ipsec dynamic set config strongcrlcheck 0
```

This setting takes effect immediately. It does not require a restart.

Enable IPSec driver event logging
To record all inbound and outbound dropped packets and other packet processing errors in the Event Viewer System log, you can set the IPSec driver event logging level to **7**. To do this, type the following at the command prompt:

```
netsh ipsec dynamic set config ipsecdiagnostics 7
```

This setting will not take effect until you restart your computer.

Notes

- If you configured an IPSec filter action on a computer to block traffic, setting the IPSec driver event logging level to **7** on that computer might generate many events that fill the System log very quickly. Accordingly, you should set the IPSec driver event logging level to **7** only for testing.
- All packets that are dropped by IPSec contribute to the **Datagrams Received Discarded** or **Datagrams Outbound Discarded** System Monitor counters for the IP object. IPSec itself does not

provide a counter object. You can monitor IPSec performance counters manually by using IP Security Monitor.

Permit inbound and outbound traffic during computer startup

To permit inbound and outbound traffic during computer startup (before the IPSec service starts), you can use the following Netsh command:

netsh ipsec dynamic set config bootmode permit

This setting will not take effect until you restart your computer.

Note

- If you start the computer in Safe Mode with Networking, the IPSec service will not start and cannot retrieve IPSec policy settings from Active Directory or the local registry, and persistent policy cannot be applied. As a result, inbound and outbound traffic will continue to be permitted until the computer is no longer in Safe Mode.

Exempt all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering

In Windows 2000 and Windows XP, by default, the IPSec driver exempts all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering. In the Windows Server 2003 family, only IKE traffic is exempt from IPSec filtering, and you can configure, block, or permit filter actions specifically for multicast and broadcast traffic (IPSec does not negotiate SAs for multicast and broadcast traffic). To restore the IPSec driver to the default Windows 2000 and Windows XP filtering behavior, you can use the following Netsh command:

netsh ipsec dynamic set config ipsecexempt 0

This setting will not take effect until you restart your computer.

Caution

- The Windows 2000 and Windows XP default exemption settings for IPSec are designed for corporate LAN environments with a low risk of attack. For this reason, you should use only the Windows 2000 and Windows XP default exemption settings when necessary for troubleshooting, in low-risk environments, or when you cannot solve program compatibility issues by configuring explicit filters in IPSec policies.

You can also exempt all broadcast, multicast, IKE, Kerberos, and RSVP traffic from IPSec filtering by modifying the registry as follows:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC\NoDefaultExempt DWORD** registry setting to a value of **0**.

2. The **NoDefaultExempt** key does not exist by default and must be created.

   Caution

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

3. Restart the computer.

Using Ipseccmd.exe to manage and monitor IPSec on computers running Windows XP

For computers running Windows XP, you can use the Ipseccmd.exe command-line tool to script IPSec policy creation and to display IPSec policy assignments, active SAs, and detailed IPSec policy settings, including filters, filter actions, and authentication methods. Ipseccmd.exe replaces and enhances the functionality provided by the Ipsecpol.exe (used in Windows 2000 for scripting IPSec policy creation) and the Netdiag.exe **/test:ipsec /v /debug** command (used in Windows 2000 to display configuration information and SAs for an active IPSec policy).

To install Ipseccmd.exe on a computer running Windows XP, run the Setup.exe file from the \Support\Tools folder on the Windows XP CD.

Notes

- Ipseccmd.exe is provided only with Windows XP, and you can use this tool to manage and monitor IPSec policies only on computers running Windows XP.
- To display all IPSec policy settings and statistics for diagnostics, use the **ipseccmd show all** command.

Using Netdiag.exe to display IPSec information and to test and view network configuration

You can use the Netdiag.exe command-line tool to display information about IPSec policies and statistics, report network configuration, and test basic networking capabilities and domain-based functionality. The following table describes how to install Netdiag.exe and provides information about functional changes relevant to IPSec for each version of this tool.

On computers running	Use this method to install	Notes
Windows Server 2003 family	For the Windows Server 2003 family, Netdiag.exe is available for installation from the Windows Server 2003 family CD (run the Suptools.msi file from the \Support\Tools folder, and choose the **Complete** setup option.	• Although you can still use Netdiag.exe to obtain basic networking information that is not IPSec-specific, for the Windows Server 2003 family, the **netdiag /test:ipsec** option has been removed, and the Netsh commands for IPSec replace all IPSec-specific functionality. To view details about IPSec policies, use either the **netsh ipsec static show** command or the **netsh ipsec dynamic show** command.
Windows XP	For Windows XP, Netdiag.exe is available for installation from the Windows XP CD (run the Setup.exe file from the \Support\Tools folder, and choose the **Complete** setup option).	• Displays information about IPSec policy assignments, including the name of the active IPSec policy, the name of the Group Policy object that assigned the policy, and the policy path. • The display of detailed IPSec policy information and IPSec statistics (provided by the **/debug** and **/v** options in Windows 2000) is not supported in Windows XP. Instead, to view details about IPSec policies, use the

		Ipseccmd.exe **show** command.
Windows 2000	For Windows 2000, an updated version of Netdiag.exe is available for download from the Web. For information, see "Windows 2000 Resource Kits," at the Microsoft Windows Resource Kits Web site.	• Netdiag.exe was enhanced in Windows 2000 Service Pack 1 to display the number of bytes offloaded by the IPSec driver to a network adapter that is capable of IPSec hardware offload. The hardware offload statistics are Offloaded Bytes Sent and Offloaded Bytes Received. • To use the **/debug** option to view details about Active Directory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory.

You can use the **netdiag /v /l** command to obtain networking information that is not IPSec-specific, for any platform. Typical uses of these options include:

- Reporting the IP configuration and routing configuration for a computer with one command.
- Testing WINS and DNS name resolution and consistency.
- Reporting the build version of a computer and the hotfixes that are installed on that computer.
- Testing the validity of domain membership, whether domain members can successfully contact domain controllers, and trust relationships.

The **/l** option generates a Netdiag.log file. This file is written to the folder in which Netdiag.exe is run.

Notes

- To run all Netdiag.exe commands, you must be a member of the Administrators group on the local computer.
- Each version of Netdiag.exe is customized to run on a different version of the Windows operating system (that is, Windows 2000, Windows XP, or the Windows Server 2003 family). You can only run Netdiag.exe on a computer that is running the Windows operating system for which Netdiag.exe has been specifically designed.
- If you are running a version of Netdiag.exe that is designed for Windows 2000 or Windows XP in a Windows Server 2003 domain, Netdiag.exe might not report the correct operating system version of the domain.
- If you run Netdiag.exe on a computer that is assigned an IPSec policy that affects most inbound and outbound IP traffic on that computer, Netdiag.exe might not report results correctly the first time it is run. This is because it might take a few seconds for IPSec to establish SAs to the many remote computers for which Netdiag.exe tests connectivity. If this occurs, wait five seconds, and then run Netdiag.exe again.

Disabling TCP/IP and IPSec hardware acceleration

Network adapters can accelerate IPSec processing by performing hardware offload of IPSec cryptographic functions (IPSec offload) and the calculation of TCP checksums (checksum offload). Such adapters might also process large TCP segments for very fast transmission (large-send offload). By default, when a network adapter driver that can perform hardware offload is enabled during the Plug and Play initialization process, the driver advertises this capability to TCP/IP and IPSec. TCP/IP and IPSec then offload tasks to the network adapter driver as appropriate.

If you use network adapters to perform hardware offload of IPSec cryptographic functions, you might need to verify that hardware acceleration is not causing problems with the packet processing that is performed by the network adapter. To do so, you can:

- Disable hardware offload functions in the network adapter driver (if this function is supported by the driver). This task does not require you to restart the computer.
- Disable TCP/IP and IPSec hardware acceleration. This task requires you to restart the computer.
- Disable only IPSec hardware acceleration. This task requires you to restart the computer.

Disabling hardware offload functions in the network adapter driver

To determine whether you can disable hardware offload functions in the network adapter driver, do one of the following:

- Open Device Manager, click **Network Adapters**, right-click the network adapter whose properties you want to view, and then click **Properties**. For more information, see Change network adapter settings
- Open Network Connections, right-click the network adapter that you want to view, click **Properties**, and then, on the **General** tab, click **Configure**.

Disabling TCP/IP and IPSec hardware acceleration

Caution

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

To disable both TCP/IP and IPSec hardware acceleration by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\DisableTaskOffload** DWORD registry setting to a value of **1**.

   Caution

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

2. Restart the computer.

If the **DisableTaskOffload** DWORD registry setting is set to a value of **0**, and if the network adapter advertises hardware offload capabilities, IPSec and TCP/IP will attempt to offload the appropriate functions to the network adapter.

Disable only IPSec hardware acceleration

To disable only IPSec hardware acceleration by modifying the registry, do the following:

1. Set the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC\EnableOffload** DWORD registry setting to a value of **0**.

   Caution

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

2. Restart the computer.

Note

- To help minimize the potential for hardware acceleration to cause problems with packet processing performed by the network adapter, make sure that you use the latest network adapter driver available from the vendor.

How to troubleshoot TCP/IP connectivity with Windows XP

TCP/IP troubleshooting tools

The following list shows some of the TCP/IP diagnostic tools that are included with Windows XP:

Basic tools

- **Network Diagnostics in Help and Support**
  Contains detailed information about the network configuration and the results of automated tests.
- **Network Connections folder**
  Contains information and configuration for all network connections on the computer. To locate the Network Connections folder, click **Start**, click **Control Panel**, and then click **Network and Internet Connections**.
- **IPConfig command**
  Displays current TCP/IP network configuration values, updates, or releases, Dynamic Host Configuration Protocol (DHCP) allocated leases, and display, register, or flush Domain Name System (DNS) names.
- **Ping command**
  Sends ICMP Echo Request messages to verify that TCP/IP is configured correctly and that a TCP/IP host is available.

Advanced tools

- **Hostname command**

  Displays the name of the host computer.
- **Nbtstat command**

  Displays the status of current NetBIOS over TCP/IP connections, updates the NetBIOS name cache, and displays the registered names and scope ID.
- **PathPing command**

  Displays a path of a TCP/IP host and packet losses at each router along the way.
- **Route command**

  Displays the IP routing table and adds or deletes IP routes.
- **Tracert command**

  Displays the path of a TCP/IP host.

To view the correct command syntax to use with each of these tools, type **-?** at a command prompt after the name of the tool.

Windows XP Professional tools

Windows XP Professional contains the following additional tools:

- **Event viewer**

  Records system errors and events.
- **Computer Management**

  Changes network interface drivers and other components.

Back to the top

Troubleshooting

The procedure that you use to troubleshoot TCP/IP issues depends on the type of network connection that you are using and the connectivity problem that you are experiencing.

Automated troubleshooting

For most issues that involve Internet connectivity, start by using the Network Diagnostics tool to identify the source of the issue. To use Network Diagnostics, follow these steps:

1. Click **Start**, and then click **Help and Support**.

2. Click the link to **Use Tools to view your computer information and diagnose problems**, and then click **Network Diagnostics** in the list on the left.
3. When you click **Scan your system**, Network Diagnostics gathers configuration information and performs automated troubleshooting of the network connection.
4. When the process is completed, look for any items that are marked "FAILED" in red, expand those categories, and then view the additional details about what the testing showed.

You can either use that information to resolve the issue or you can provide the information to a network support professional for help. If you compare the tests that failed with the documentation in the Manual Troubleshooting section later in this article, you may be able to determine the source of the issue. To interpret the results for TCP/IP, expand the Network Adapters section of the results, and then expand the network adapter that failed the testing.

You can also start the Network Diagnostics interface directly by using the following command:
netsh diag gui

Manual troubleshooting

To manually troubleshoot your TCP/IP connectivity, use the following methods in the order that they appear:

Method 1: Use the IPConfig tool to verify the configuration

To use the IPConfig tool to verify the TCP/IP configuration on the computer that is experiencing the problem, click **Start**, click **Run**, and then type **cmd**. You can now use the **ipconfig** command to determine the host computer configuration information, including the IP address, the subnet mask, and the default gateway.

The **/all** parameter for IPConfig generates a detailed configuration report for all interfaces, including any remote access adapters. You can redirect IPConfig output to a file to paste the output into other documents. To do this, type:
ipconfig > \*folder_name*\*file_name*

The output receives the specified file name and is stored in the specified folder.

You can review the IPConfig output to identify issues that exist in the computer network configuration. For example, if a computer is manually configured with an IP address that duplicates an existing IP address that is already detected, the subnet mask appears as 0.0.0.0.

If your local IP address is returned as 169.254.*y*.*z* with a subnet mask of 255.255.0.0, the IP address was assigned by the Automatic Private IP Addressing (APIPA) feature of Windows XP Professional. This assignment means that TCP/IP is configured for automatic configuration, that no DHCP server was found, and that no alternative configuration is specified. This configuration has no default gateway for the interface.

If your local IP address is returned as 0.0.0.0, the DHCP Media Sensing feature override turned on because the network adapter detected its lack of connection to a network, or TCP/IP detected an IP address that duplicates a manually configured IP address.

If you do not identify any issues in the TCP/IP configuration, go to Method 2.

Method 2: Use the Ping tool to test your connectivity

If you do not identify any issues in the TCP/IP configuration, determine whether the computer can connect to other host computers on the TCP/IP network. To do this, use the Ping tool.

The Ping tool helps you verify IP-level connectivity. The **ping** command sends an ICMP Echo Request message to a destination host. Use Ping whenever you want to verify that a host computer can send IP packets to a destination host. You can also use Ping to isolate network hardware problems and incompatible configurations.

Note If you ran the **ipconfig /all** command, and the IP configuration appeared, you do not have to ping the loopback address and your own IP address. IPConfig has already performed these tasks to display the configuration. When you troubleshoot, verify that a route exists between the local computer and a network host. To do this, use the following command:
ping *IP address*

Note*IP address* is the IP address of the network host that you want to connect to.

To use the **ping** command, follow these steps:

1. Ping the loopback address to verify that TCP/IP is installed and correctly configured on the local computer. To do this, type the following command:

   ping 127.0.0.1

   If the loopback test fails, the IP stack is not responding. This problem may occur if any one or more of the following conditions is true:

   - The TCP drivers are corrupted.
   - The network adapter is not working.
   - Another service is interfering with IP.

2. Ping the IP address of the local computer to verify that the computer was correctly added to the network. If the routing table is correct, this procedure just forwards the packet to the loopback address of 127.0.0.1. To do this, type the following command:

   ping *IP address of local host*

   If the loopback test succeeds but you cannot ping the local IP address, there may be an issue with the routing table or with the network adapter driver.

3. Ping the IP address of the default gateway to verify that the default gateway is working and that you can communicate with a local host on the local network. To do this, type the following command:

   ping *IP address of default gateway*

   If the ping fails, you may have an issue with the network adapter, the router or gateway device, the cabling, or other connectivity hardware.

4. Ping the IP address of a remote host to verify that you can communicate through a router. To do this, type the following command:

   ping *IP address of remote host*

   If the ping fails, the remote host may not be responding, or there may be a problem with the network hardware between computers. To rule out an unresponsive remote host, use Ping again to a different remote host.

5. Ping the host name of a remote host to verify that you can resolve a remote host name. To do this, type the following command:

ping ***Host name of remote host***

Ping uses name resolution to resolve a computer name into an IP address. Therefore, if you successfully ping an IP address but you cannot ping a computer name, there is a problem with host name resolution, not with network connectivity. Verify that DNS server addresses are configured for the computer, either manually in the properties of TCP/IP, or by automatic assignment. If DNS server addresses are listed when you type the **ipconfig /all** command, try to ping the server addresses to make sure that they are accessible.

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

- Make sure that the local computer's IP address is valid and that it is correct on the **General** tab of the **Internet Protocol (TCP/IP) Properties** dialog box or when it is used with the Ipconfig tool.
- Make sure that a default gateway is configured and that the link between the host and the default gateway is working. For troubleshooting purposes, make sure that only one default gateway is configured. Although you can configure more than one default gateway, gateways after the first gateway are used only if the IP stack determines that the original gateway is not working. The purpose of troubleshooting is to determine the status of the first configured gateway. Therefore, you can delete all the other gateways to simplify your task.
- Make sure that Internet Protocol security (IPSec) is not turned on. Depending on the IPSec policy, Ping packets may be blocked or may require security. For more information about IPSec, go to Method 7: Verify Internet Protocol security (IPSec).

Important If the remote computer that you are pinging is across a high-delay link such as a satellite link, response may take longer. You can use the **-w** (wait) parameter to specify a longer timeout period than the default timeout of four seconds.

Method 3: Use the PathPing tool to verify a route

The PathPing tool detects packet loss over multiple-hop paths. Run a PathPing analysis to a remote host to verify that the routers on the way to the destination are operating correctly. To do this, type the following command:
pathping ***IP address of remote host***

Method 4: Use the Arp tool to clear the ARP cache

If you can ping both the loopback address (127.0.0.1) and your IP address but you cannot ping any other IP addresses, use the Arp tool to clear out the Address Resolution Protocol (ARP) cache. To view the cache entries, type any one of the following commands:

```
arp -a
```

```
arp -g
```

To delete the entries, type the following command:

```
arp -d IP address
```

To flush the ARP cache, type the following command:

```
netsh interface ip delete arpcache
```

Method 5: Verify the default gateway

The gateway address must be on the same network as the local host. Otherwise, messages from the host computer cannot be forwarded outside the local network. If the gateway address is on the same network as the local host, make sure that the default gateway address is correct. Make sure that the default gateway is a router, not just a host. And make sure that the router is enabled to forward IP datagrams.

Method 6: Use the Tracert tool or the Route tool to verify communications

If the default gateway responds correctly, ping a remote host to make sure that network-to-network communications are working correctly. If communications are not working correctly, use the Tracert tool to trace the path of the destination. For IP routers that are Microsoft Windows 2000-based or Microsoft Windows NT 4.0-based computers, use the Route tool or the Routing and Remote Access snap-in to view the IP routing table. For other IP routers, use the vendor-designated appropriate tool or facility to examine the IP routing table.

Most frequently, you receive the following four error messages when you use Ping during troubleshooting:

TTL Expired in Transit

This error message means that the number of required hops exceeds the Time to Live (TTL). To increase TTL, by use the **ping -i** command. A routing loop may exist. Use the **Tracert** command to determine whether misconfigured routers have caused a routing loop.

Destination Host Unreachable

This error message means that no local or remote route exists for a destination host at the sending host or at a router. Troubleshoot the local host or the router's routing table.

Request Timed Out

This error message means that the Echo Reply messages were not received in the designated timeout period. By default, the designated timeout period is four seconds. Use the **ping -w** command to increase the timeout.

Ping request could not find host

This error message means that the destination host name cannot be resolved. Verify the name and the availability of DNS or WINS servers.

Method 7: Verify Internet Protocol security (Ipsec)

IPSec can improve security on a network, but changing network configurations or troubleshooting problems more difficult. Sometimes, IPSec policies require secured communication on a Windows XP Professional-based computer. These requirements can make it difficult to connect to a remote host. If IPSec is implemented locally, you can turn off the IPSEC Services service in the Services snap-in.

If the difficulties end when you stop the IPSec services, IPSec policies are either blocking the traffic or requiring security for the traffic. Ask the security administrator to modify the IPSec policy.

Method 8: Verify packet filtering

Because of mistakes in packet filtering, address resolution or connectivity may not work. To determine whether packet filtering is the source of a network problem, turn off TCP/IP packet filtering. To do this, follow these steps:

1. Click **Start**, click **Control Panel**, click **Network and Internet Connections**, and then click **Network Connections**.
2. Right-click the local area connection that you want to modify, and then click **Properties**.
3. On the **General** tab, in the **This connection uses the following items** list, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. Click **Advanced**, and then click the **Options** tab.
5. In the **Optional Settings** dialog box, click **TCP/IP Filtering**, and then click the **Properties** tab.
6. Click to clear the **Enable TCP/IP Filtering (All adapters)** check box, and then click **OK**.

To ping an address, use its DNS name, its NetBIOS computer name, or its IP address. If the ping succeeds, the packet filtering options may be misconfigured or too restrictive. For example, the filtering can allow the computer to act as a Web server, but, to do this, the filtering may turn off tools such as remote administration. To restore a wider range of permissible filtering options, change the permitted values for the TCP port, the UDP port, and the IP protocol.

Method 9: Verify the connection to a specific server

To determine the cause of connectivity problems when you are trying to connect to a specific server through NetBIOS-based connections, use the **nbtstat -n** command on the server to determine what name the server registered on the network.

The **nbtstat -n output** command lists several names that the computer has registered. The list will include a name that looks similar to the computer's name that is configured on the **Computer Name** tab under **System** in Control Panel. If not, try one of the other unique names that the **nbtstat** command displays.

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

Method 10: Verify remote connections

To determine why a TCP/IP connection to a remote computer stops responding, use the **netstat -a** command to show the status of all activity for TCP and UDP ports on the local computer.

Typically, a good TCP connection shows 0 bytes in the **Sent** and **Received** queues. If data is blocked in either queue or the state of the queues is irregular, the connection may be faulty. If data is not blocked, and the state of the queues is typical, you may be experiencing network or program delay.

Method 11: Use the Route tool to examine the routing table

For two hosts to exchange IP datagrams, both hosts must have a route to each other, or they must use default gateways that have a route. To view the routing table on a Windows XP-based host, type the following command:

route print

Method 12: Use the Tracert tool to examine paths

Tracert sends ICMP Echo Request messages that have incrementally higher values in the IP header TTL field to determine the path from one host to another through a network. Then Tracert analyzes the ICMP messages that are returned. With Tracert, you can track the path from router to router for up to 30 hops. If a router has failed, or the packet is routed into a loop, Tracert reveals the problem. After you locate the problem router, you can contact the router administrator if the router is offsite, or you can restore the router to fully functional status if the router is under your control.

Method 13: Troubleshoot gateways

If you receive the following error message during configuration, determine whether the default gateway is located on the same logical network as the computer's network adapter:

Your default gateway does not belong to one of the configured interfaces

Compare the network ID part of the default gateway IP address with the network IDs of the computer's network adapters. Specifically, verify that the bitwise logical **AND** of the IP address and the subnet mask equals the bitwise logical **AND** of the default gateway and the subnet mask.

For example, a computer that has a single network adapter that is configured with an IP address of 172.16.27.139 and a subnet mask of 255.255.0.0 must use a default gateway of the form 172.16.*y*.*z*. The network ID for this IP interface is 172.16.0.0.

Additional resources

The following resources contain additional information about how to troubleshoot Microsoft TCP/IP:

See the "Configuring TCP/IP" topic in the documentation for the Microsoft Windows XP Professional Resource Kit.

See "Introduction to TCP/IP" in the *TCP/IP Core Networking Guide* of the Microsoft Windows 2000 Server Resource Kit for general information about the TCP/IP protocol suite.

See "Unicast Routing Overview" in the *Internetworking Guide* of the Microsoft Windows 2000 Server Resource Kit for more information about routing principles.

See "TCP/IP Troubleshooting" in the *TCP/IP Core Networking Guide* of the Microsoft Windows 2000 Server Resource Kit for more information about IP packet filtering.
BECSN - Building Enhanced Cisco Security Networks v2.0

Managing IGMP traffic with Cisco NetFlow

With Cisco NetFlow technology and Plixer’s Scrutinizer NetFlow Analyzer and Flow Analytics module, network administrators can now monitor and alert on unwanted transport protocols, such as IGMP.

In the Flow Analytics gadget displayed below — the Top Transport gadget — four transport protocols are listed. The three listed in white are in the “Allowed transports” list as shown in the lower half of the image, whereas the IGMP protocol is highlighted in yellow, indicating a violation of the Top Network Transports algorithm. To display and allow editing of the “Allowed transports” list, click on Settings at the bottom of the Top Transport gadget.

To add this protocol to the “Allowed list”, click on the red plus (+) sign to the right of the protocol entry. This will stop this protocol from violating the Top Network Transports algorithm.

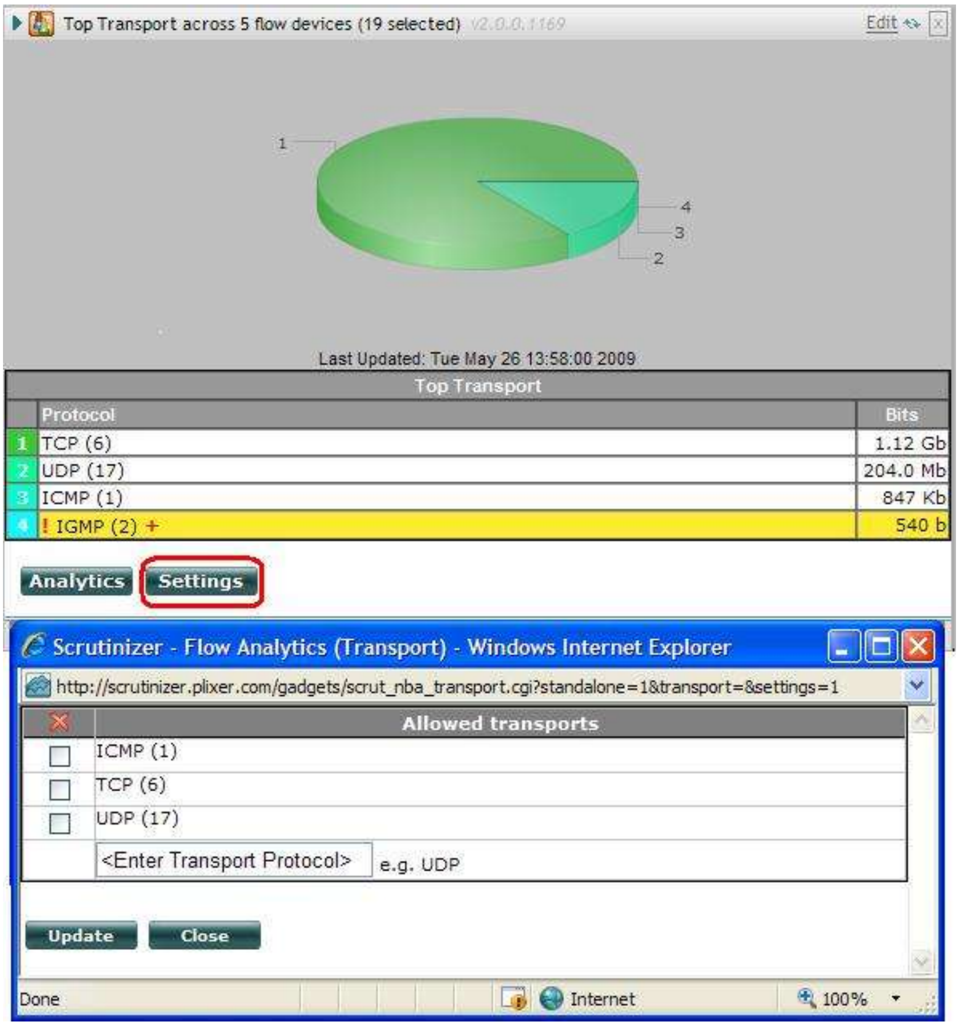

Top Transport across 5 flow devices (19 selected) v2.0.0.1169

Last Updated: Tue May 26 13:58:00 2009

Top Transport

	Protocol	Bits
1	TCP (6)	1.12 Gb
2	UDP (17)	204.0 Mb
3	ICMP (1)	847 Kb
4	! IGMP (2) +	540 b

Scrutinizer - Flow Analytics (Transport) - Windows Internet Explorer

http://scrutinizer.plixer.com/gadgets/scrut_nba_transport.cgi?standalone=1&transport=&settings=1

	Allowed transports
□	ICMP (1)
□	TCP (6)
□	UDP (17)
	<Enter Transport Protocol> e.g. UDP

Clicking on the red exclamation point (!) to the left of the protocol’s entry will open a new Alarms window showing the alarms for IGMP.

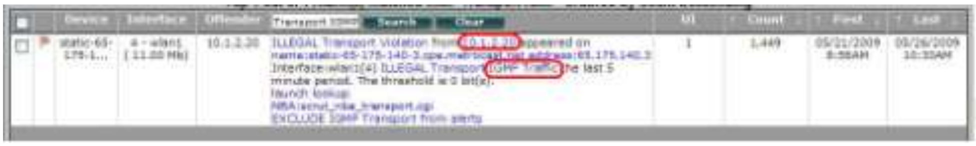

□		Device	Interface	Offender	Transport IGMP Search Clear	UI	Count	First	Last
□	⚑	static-65-175-1...	4 - wlan1 (11.00 Mb)	10.1.2.20	ILLEGAL Transport Violation from 10.1.2.20 appeared on name:static-65-175-140-3.cpe.metrocast.net address:65.175.140.3 Interface:wlan1(4) ILLEGAL Transport IGMP Traffic the last 5 minute period. The threshold is 0 bit(s). launch lookup NBA:scrut_nba_transport.cgi EXCLUDE IGMP Transport from alerts	1	1,449	05/21/2009 8:58AM	05/26/2009 10:33AM

From this Alarms page, you can exclude the violating host (10.1.2.20), by clicking on 10.1.2.20 in the message section of the alarm.

Hovering over the text “ILLEGAL Transport IGMP Traffic” will display how much traffic has been transmitted for this protocol from this IP address.

To receive email alerts based on this illegal transport traffic, the Top Network Transports algorithm in Flow Analytics can be configured to send syslogs to your syslog server (Logalot can be used here), with the syslog server generating email alerts.

Another handy tool from Plixer International and Scrutinizer NetFlow Analyzer.

Windows Internet Name Service

From Wikipedia, the free encyclopedia
Jump to: navigation, search

Windows Internet Name Service (WINS) is Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names. Effectively WINS is to NetBIOS names, what DNS is to domain names — a central mapping of host names to network addresses. Like DNS it is broken into two parts, a Server Service (that manages the encoded Jet Database, server to server replication, service requests, and conflicts) and a TCP/IP Client component which manages the client's registration and renewal of names, and takes care of queries.

[edit] Overview

Network address mappings are dynamically updated so that when a client needs to contact another computer on the network it can get its up-to-date IP address which may be issued by a DHCP server. Aside from this the WINS functionality provides a way of keeping the client names unique on the network.

Networks normally have more than one WINS server and each WINS server should be in push/pull replication; where more than two WINS servers are required the best practice replication model is the hub and spoke, thus the WINS design is not central but distributed. Each WINS server holds a full copy of every other related WINS system's records. There is no hierarchy in WINS (unlike DNS), but like DNS its database can be queried for the address to contact rather than broadcasting a request for which address to contact. The system therefore reduces broadcast traffic on the network, however replication traffic can add to WAN/LAN traffic, although this can be set to replicate in non busy periods. By design any WINS client can register any name with any WINS server. This makes the system prone to abuse or unreliable through poor administration.

All WINS clients should be configured to use a primary WINS server and a different secondary WINS server. The secondary would normally be the hub server. The setting of which WINS servers to use is either in the DHCP scope options or a per client hard coded value.

As of Windows 2000, DNS provides the favored alternative to WINS, as part of Active Directory.[1]

In theory, if DNS is available, WINS is only necessary if pre-Windows 2000 clients or servers need to resolve names. In reality, especially in large enterprise environments, applications such as SMS 2003 with its use of the 1A record, MS SQL Server 2000 for use of named pipes, and Exchange Server 2000 and 2003 both require WINS for full functionality.[2]

The WINS server from Microsoft is only available as a service to run on the Windows Server family of operating systems. The WINS client from Microsoft is common across all its operating systems including DOS. WINS clients can also be devices such as IP phones and printers.

[edit] References

1. ^ "What is WINS? - WINS and DNS". Microsoft. 2003-03-28. http://technet2.microsoft.com/WindowsServer/en/library/01b8e158-4587-4269-917a-5ad38c2537021033.mspx?mfr=true. Retrieved 2007-03-19.
2. ^ "Exchange Server 2003 and Exchange 2000 Server require NetBIOS name resolution for full functionality". Microsoft. 2007-10-25. http://support.microsoft.com/kb/837391.

[edit] External links

- Official sources:
    - Microsoft Technet Docs.
    - Microsoft Technet: WINS on Windows Server 2003
    - MSKB: 837391: Exchange Server 2003 and Exchange 2000 Server require NetBIOS name resolution for full functionality
- Other:
    - Name Resolution chapter in Using Samba online book (also published by O'Reilly as ISBN 0-596-00256-4), which speaks about WINS.

Retrieved from "http://en.wikipedia.org/wiki/Windows_Internet_Name_Service"

Domain Name System

From Wikipedia, the free encyclopedia
Jump to: navigation, search

The **Domain Name System** (**DNS**) is a hierarchical naming system built on a distributed database for computers, services, or any resource connected to the Internet or a private network. It associates various information with domain names assigned to each of the participating entities. Most importantly, it translates domain names meaningful to humans into the numerical identifiers associated with networking equipment for the purpose of locating and addressing these devices worldwide.

An often-used analogy to explain the Domain Name System is that it serves as the *phone book* for the Internet by translating human-friendly computer hostnames into IP addresses. For example, the domain name *www.example.com* translates to the addresses *192.0.32.10* (IPv4) and *2620:0:2d0:200::10* (IPv6).

The Domain Name System makes

Windows 2000 Server Windows Internet Naming Service (WINS) Overview

Published: April 19, 1999

Abstract

The Microsoft Windows 2000 operating system Internet Name Service (WINS) introduces new features and enhancements to the WINS server, WINS client, and WINS Manager. WINS provides a distributed database for registering and querying dynamic computer name-to-IP address mapping in a routed network environment.

On This Page

Introduction

Microsoft Windows Internet Name Service (WINS) has been enhanced for the release of Microsoft Windows 2000 Server. The result is an easier-to-manage and more robust solution for mapping NetBIOS names to IP addresses on Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

Windows 2000 WINS includes server enhancements, additional client functions, and an improved management tool. WINS provides a distributed database for registering and querying dynamic computer name-to-IP address mapping in a routed network environment. This support for dynamic registering of NetBIOS computer names means that WINS can be used with Dynamic Host Configuration Protocol (DHCP) services to provide easy configuration and administration of Windows-based TCP/IP networks.

The WINS server solves the problems inherent in resolving NetBIOS names through IP broadcasts, and frees network administrators from the demands of updating static mapping files, such as LMHOST files. WINS, which is compliant with the NetBIOS Name Server (NBNS) RFCs (1001/1002), also automatically updates the WINS database when dynamic addressing through DHCP results in new IP addresses for computers that move between subnets. Neither the user nor the network administrator needs to make manual accommodations for such name resolutions.

The new implementation of WINS provides a number of features, including:

- **Persistent connections**—This configurable feature allows each WINS server to maintain a persistent connection with one or more replication partners to eliminate the overhead of opening and terminating connections and to increase the speed of replication.
- **Manual tombstoning**—Use of the Manual tombstoning feature marks a record for deletion so that the tombstone state for the record is replicated across all WINS servers, preventing an undeleted copy of the record on a different server database from being re-propagated.
- **Improved management tools**—The WINS Manager is fully integrated with the Microsoft Management Console (MMC), providing a more user-friendly and powerful environment for viewing and managing WINS information.
- **Enhanced filtering and record searching**—These functions help locate records of interest by showing only those that fit a specific criteria. This is particularly useful for analyzing very large WINS databases.
- **Dynamic record deletion and multi**-select—Managing the WINS database is made easier with dynamic record deletion and multi-select. Dynamic and static records can be deleted, and the point-and-click interface makes it possible to delete files with non-alphanumeric characters that could not be handled from the command line.
- **Record verification and version number validation**—Two tools are available for quickly checking the consistency between various WINS servers. The tests are done by comparing the IP addresses of a NetBIOS name query returned from different WINS servers or by examining owner address to version-number mapping tables.

- **Export function**—The Export command can be used to place WINS data into a comma-delimited text file that can be imported into Microsoft Excel, reporting tools, scripting applications, and so on, for analysis and reporting.
- **Increased fault tolerance**—Windows 2000 and Windows 98 allow a client to specify more than two WINS servers (up to a maximum of 12 addresses) per interface. The extra WINS server addresses are used only if the primary and secondary WINS servers fail to respond.
- **Dynamic re-registration**—WINS clients can now re-register their NetBIOS name-to- IP address mapping without rebooting the server.

All of this combines to make Windows 2000 WINS a superior choice for NetBIOS name resolution. The new generation of WINS services are designed to make many network management tasks much easier for network managers.

Top of page

WINS Functional Description

Please refer to the white paper "Microsoft Windows NT Server 4.0 Windows Internet Naming Service (WINS) Architecture and Capacity Planning" (available at http://support.microsoft.com/default.aspx?scid=kb;en-us;239950) for a detailed description of WINS elements and functions.

Top of page

New Features Of Windows 2000 WINS

Windows 2000 WINS contains significant enhancements, many of which were suggested by network managers. The result is an even more powerful and easier-to-manage NetBIOS-to-IP address service. New features include:

- Persistent Connections
- Manual Tombstoning
- Improved Management Tool
- Enhanced Filtering and Record Searching
- Dynamic Record Deletion and Multi-select
- Record Verification and Version Number Validation
- Export Function
- Increased Fault Tolerance
- Dynamic Re-registration

Persistent Connections

Windows 2000 WINS introduces persistent connections between WINS server replication partners. This is important because the WINS database is collectively managed by a set of WINS servers, each of which has a copy of the WINS database. To keep these copies consistent, servers replicate records among each other. Each WINS server is configured with a set of one or more replication partners. Each new computer added to or substituted on the network registers its name and IP address with a WINS server, which in turn propagates the new record to all other WINS servers in the enterprise. The result is that every server has the record pertaining to that new computer.

Earlier versions of WINS required a new connection to be established between the WINS servers whenever replication was to occur, each of which required a modest number of processor cycles.

Therefore, network managers would set their systems to accumulate a configurable number of records prior to having servers establish connection with replication partners. This caused a delay to be introduced in the updating of the entire database, perhaps as long as several minutes, which could cause windows of inconsistency with replication partners.

Windows 2000 WINS provides a configurable feature that allows a server to request a persistent connection with one or more replication partners, which eliminates the overhead of opening and terminating connections. Persistent connections increase the speed of replication because a server can immediately send records to its partners without incurring the cost of establishing temporary connections each time. This provides the opportunity for every record received be immediately updated across the network, making records more consistent. The bandwidth used by persistent connections is minimal because the connection is usually idle.

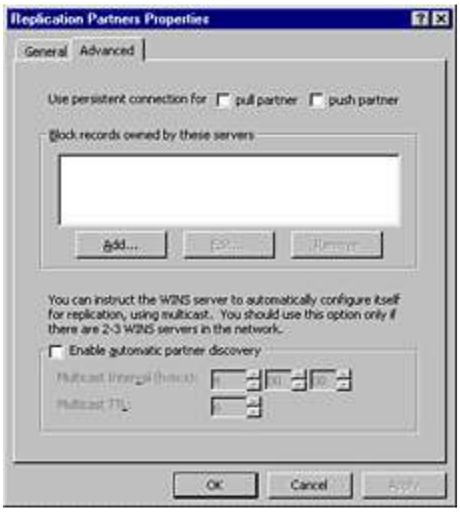

Replication Partners Properties

Use persistent connection for □ pull partner □ push partner

Block records owned by these servers

You can instruct the WINS server to automatically configure itself for replication, using multicast. You should use this option only if there are 2-3 WINS servers in the network.

□ Enable automatic partner discovery

Figure 1: Windows 2000 WINS configured to Use persistent connection

Manual Tombstoning

Windows 2000 WINS gives network managers the ability to manually tombstone records, marking them for deletion so that the order to delete is propagated across all WINS servers. This is significant because WINS is a distributed environment. Each server holds the entire database, and replication is used to propagate updates to every server.

In earlier versions of WINS, the removal of unwanted records could be difficult. Records were marked for removal on only one server, and that information is replicated to other WINS replication partners. Depending upon replication configuration, record removal might not occur correctly.

The manual tombstoning option of Windows 2000 WINS addresses such problems. The timing of the tombstoned state exceeds the propagation delay incurred with replication across the network. When the time limit is reached, tombstoned records are removed on all servers. Manual tombstoning provides an excellent way of dealing with static records, too. Manual tombstoning is available from both the WINS graphical user interface and the WINS command-line interface.

In WINS Manager, tombstoning is displayed as an option under the record deletion function. A record can be either deleted or tombstoned. The ability to manually tombstone records requires Windows 2000–enabled WINS servers, but tombstoned records replicate to Windows NT 3.51 and Windows NT 4.0 WINS servers.

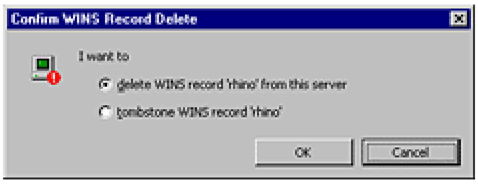

Confirm WINS Record Delete

I want to

● delete WINS record 'rhino' from this server

○ tombstone WINS record 'rhino'

Figure 2: Manual tombstoning function in the WINS Manager

For a more information on tombstoning, see the white paper, “Microsoft Windows NT Server 4.0 Windows Internet Naming Service (WINS) Architecture and Capacity Planning.”

Improved Management Tool

The Windows 2000 WINS Manager has an updated graphical user interface, incorporating all of the user-friendly features of the Windows interface, including resizable windows.

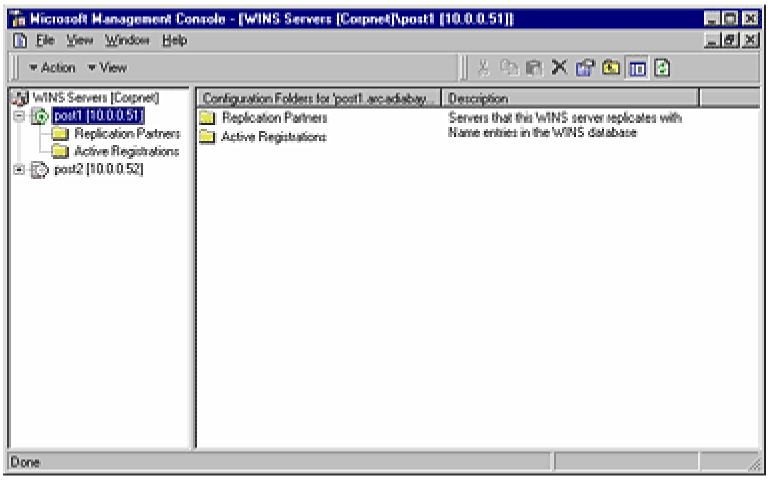

Configuration Folders for 'post1.arcadiabay...	Description
Replication Partners	Servers that this WINS server replicates with
Active Registrations	Name entries in the WINS database

Figure 3: The Microsoft Management Console user-friendly interface

Microsoft Management Console (MMC)

The Windows 2000 WINS Manager is fully integrated with the robust Microsoft Management Console (MMC). This provides network managers with the consistency of using the MMC for all of their administrative tasks. With MMC, multiple WINS Manager windows can be opened at once.

Higher Performance

A major enhancement to the Windows 2000 WINS Manager is a multithreaded user interface (UI), which allows background tasks to take place while a foreground task is being performed, providing much faster response. Multithreading of the UI permits a network manager to run multiple UI tasks simultaneously. A manager can select the Active Registration node to display the database, and then—rather than waiting for the complete database to be displayed—can select the node of another server, change the configuration of replication partners, or perform other tasks.

Resizable Columns

The window for database records now features resizable display columns. Using standard left-right drag, users can resize columns to accommodate their current tasks.

List View

List View provides a flexible, Explorer-like, function for sorting information according to column type. This means network managers can, with the click of the mouse, sort data according to:

- Record name
- IP address
- Type
- Static/dynamic
- Status
- Expiration date

- Version number

Enhanced Filtering and Record Searching

A Quick Find function has been added to the WINS Manager to help network managers and support staff search for only a few records without downloading and viewing the entire database.

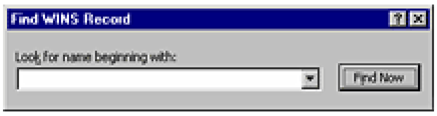

Figure 4: The new Quick Find function

For example, a network manager may want to review all WINS records that contain the location name for St. Louis, MO. This would allow the network manager to selectively view all records that contain the STL string, without requiring a retrieval of the entire database. The Quick Find function selects the records containing the STL string and presents only those records to the network manager.

Filtering capability eliminates the need to examine every record in an entire database by stating which type of record is desired. After an entire database is downloaded, a search can be refined to a certain type of computer. (For more detail on NetBIOS name registrations and their types, see the white paper, "Microsoft Windows NT Server 4.0 Windows Internet Naming Service (WINS) Architecture and Capacity Planning.") Searches can also be done for domain controllers only, and multiple filters can be applied to view a combined display of records for computers that act as RAS servers, file servers, or other types of computers, and in various combinations.

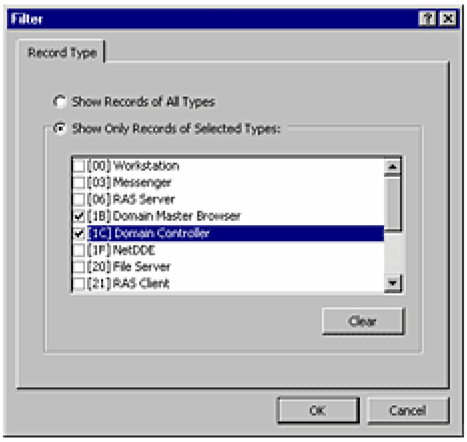

○ Show Records of All Types
● Show Only Records of Selected Types:

□ [00] Workstation
□ [03] Messenger
□ [06] RAS Server
☑ [1B] Domain Master Browser
☑ [1C] Domain Controller
□ [1F] NetDDE
□ [20] File Server
□ [21] RAS Client

Figure 5: Searching by computer type

Dynamic Record Deletion and Multi-Select

In addition to static record deletion available in previous versions of WINS, dynamic records can now be deleted. Multi-select can be used to delete a number of records at once. To select a range of contiguous records, click the first record to delete, and then SHIFT+click the last record in the range to delete. This delete function also includes manual tombstoning.

The point-and-click record deletion is especially helpful when dealing with file names containing non-alphanumeric characters, which prevent them from being deleted from the command line.

Record Verification and Version Number Validation

Record verification allows network managers to verify a set of WINS records against a set of WINS servers by sending name queries to each server to ensure that the names are consistent among all servers.

Version number validation ensures consistency of version IDs between servers. It first obtains the owner address-to- version number maps from different WINS servers. It then checks the consistency of their databases by ensuring that a WINS server always has the highest version number among the network of WINS servers for records owned by it.

Consistency Checking

Consistency checking can be done from the MMC. Network managers can configure the frequency with which the system compares the WINS databases. However, consistency checking is very network bandwidth-intensive and consumes a great many cycles on the WINS server. To perform consistency checking, a WINS replicates all of a particular owner's records from another WINS, and then checks them to determine whether its database is synchronized with the other WINS for that owner.

Autodiscovery of WINS Partners

The autodiscovery feature, which can be turned on from the MMC, enables a WINS server to automatically discover its replication partners. Periodically WINS servers announce their presence on the network. WINS servers that have the autodiscovery feature turned on listen for these announcements and learn about other WINS servers on the network. WINS servers discovered in this way are automatically added to the partners list as both a push and pull partner. The WINS announcements are sent by multicast on a multicast address reserved for WINS. (224.0.1.24). This feature should not be used if there are more than three WINS servers on the network.

Monitoring

The monitoring features of MMC provide network managers with the status of WINS servers. Monitoring can check for consistency between servers, detect version-number inconsistencies, detect communication failures between WINS servers, and verify replication configuration setup.

Export Function

Windows 2000 WINS Manager provides an Export command,. To see this command, right-click the Active Registration folder. The Export function copies any data on the right-hand side of the MMC window into a comma-delimited text file. The exported data can then be imported into Microsoft Excel, reporting tools, scripting applications, or other files. This function permits network managers use whatever tools they choose for generating reports, analyzing network performance, and other functions.

Increased Fault Tolerance

Windows 2000 and Windows 98 provide an extra measure of fault tolerance by allowing a client to specify more than two WINS servers (up to a maximum of 12 addresses) per interface through either DHCP or Setup. The extra WINS server addresses are used to resolve names only if the primary and secondary WINS servers fail to respond. If one of the extra WINS servers is used, its name resolutions are cached to be used the next time the primary and secondary WINS servers fail to resolve the name.

Dynamic Re-registration

WINS clients have been improved to allow them to re-register their NetBIOS name-to-IP address mappings without rebooting the computers. This is useful when incorrect static entries exists or if a WINS database is restored with an old record, for example. The version ID is updated on the server to cause re-replication.

Top of page

Summary

Microsoft Windows Internet Name for Windows 2000 Server provides a solid platform for managing NetBIOS name-to-IP address resolution on a routed TCP/IP network. The enhancements and new features in this version of WINS are to a great extent influenced by feedback from network managers in the field.

Persistent connections allow each WINS server to maintain a connection with one or more replication partners to eliminate the overhead of opening and terminating connections and to increase the speed of replication.

Manual tombstoning ensures that records marked for deletion are deleted across the distributed database to prevent records from re-propagating back across the WINS servers.

An improved WINS Manager tool, including full integration with the Microsoft Management Console, provides a more user-friendly and powerful environment for viewing and managing WINS information. In addition, the searching capability speeds up database record retrieval, and filtering cuts down on the number of records that need to be displayed.

Managing the WINS database is made easier with dynamic record deletion and multi-select. Dynamic and static records can be deleted, and the point-and-click interface makes it possible to delete files with non-alphanumeric characters that could not be handled from the command line.

Record verification and version-number validation provide quick checks for consistency between WINS servers. Full database consistency checking is available right from the WINS Manager.

The Export command can be used to place WINS data into a comma-delimited text file that can be imported into Microsoft Excel, reporting tools, scripting applications, and so on, for analysis and reporting.

WINS clients can now re-register without rebooting the server, and they have better fault tolerance because they can query against more than two WINS servers.

These features combine to make the Windows 2000 Windows Internet Name Service a much improved solution for managing dynamic NetBIOS name resolution across a routed TCP/IP network.

Configuring and Troubleshooting TCP/IP

TCP/IP encompasses a vast array of utilities and network services. This suite of services has evolved to become the industry standard for both the Internet and for local area networks (LANs) using personal computer network operating systems like Novell NetWare 5, Unix, and Windows XP.

TCP/IP is the default protocol when you install Windows XP Professional. It provides a means for connecting dissimilar computer systems. TCP/IP scales well and is typically the best choice for any size of organization. TCP/IP and its name resolution partner, Domain Name System (DNS), are both required components for implementing Active Directory in the Windows 2000/2003 Server family of products.

Deciphering the TCP/IP Protocol Suite for Windows XP

TCP/IP is more than a standardized specification for data transport over a network wire. It is a sophisticated toolbox of data transport services, name resolution services, and troubleshooting utilities. Microsoft's implementation of TCP/IP for Windows XP includes the following network services and components:

- Dynamic Host Configuration Protocol (DHCP) This service is based on an industry-standard specification for automatically assigning (or leasing) IP addresses to computers connected to the network. The addresses are assigned from a predefined pool (or scope) of IP addresses that an administrator must configure. DHCP makes the chore of assigning and maintaining TCP/IP addresses on hundreds or thousands of computers much easier than having to maintain an exhaustive list of IP addresses and computer names by hand. However, administrators should manually assign static IP addresses for domain controllers, file and print server computers, and printers. You can install the DHCP service only in the Windows 2000/2003 Server product line, but DHCP can assign addresses to both servers and workstations. Any operating system that can make DHCP-enabled requests for IP addresses can use a DHCP server that is running

Windows 2000. DHCP-enabled operating systems include Windows 3.x, 9x, ME, NT, 2000, 2003 and XP.

- DNS server Computers understand and work well with numbers, but humans remember names much more easily than numbers. TCP/IP requires that each network device be assigned a numeric IP address. DNS, in conjunction with DNS servers, maps numeric IP addresses to computer (host) names and vice versa. DNS employs a hierarchical system of domains and subdomains that helps to make this name resolution service very scalable. DNS servers mitigate the need for a manually maintained HOSTS file to be stored on each computer. Windows 2000/2003 DNS servers offer added functionality such as Active Directory Integrated Zones, Incremental Zone Transfers, and Secure Dynamic Updates. DNS is a requirement for implementing Active Directory.
- Windows Internet Naming Service (WINS) This service is Microsoft's implementation of a name resolution mechanism to match IP addresses to NetBIOS computer names and vice versa. WINS servers can greatly reduce NetBIOS traffic on networks by decreasing the amount of broadcast traffic that occurs when computers attempt to resolve unknown NetBIOS computer names to IP addresses. For an Active Directorybased network in Windows 2000/2003 native mode with no applications that require NetBIOS, nor any legacy Windows clients, WINS becomes unnecessary.
- Automatic Private IP Addressing (APIPA) Microsoft first introduced this feature in Windows 98. For computers that are configured to obtain an IP address automatically, APIPA kicks in if no DHCP server is available on the network to lease out an IP address. APIPA automatically queries the other computers on the network to ensure it does not duplicate an IP address, and then assigns a unique IP address to the local computer using the IP address scheme of 169.254.x.y with the subnet mask of 255.255.0.0. The Internet Assigned Numbers Authority (IANA) has reserved the IP address range of 169.254.0.0 through 169.254.255.255 for APIPA. This ensures that any IP address that APIPA generates does not conflict with any public, routable addresses. This feature is turned on by default in Windows XP Professional.
- Serial Line Internet Protocol (SLIP) This specification is an older Unix standard for serial communications. Windows XP supports SLIP for backward-compatibility purposes. You can use SLIP only for outbound connections on Windows XP Professional.
- Point-to-Point Protocol (PPP) PPP has effectively replaced SLIP. PPP is a remote access/dial-up protocol that supports industry-standard network protocols such as TCP/IP, NWLink, NetBEUI, and AppleTalk. PPP is optimized for low-bandwidth connections, so it is the preferred remote access protocol for dial-up/modem connections.
- Point-to-Point Tunneling Protocol (PPTP) The only Virtual Private Network (VPN) protocol that shipped with Windows NT 4, PPTP encapsulates TCP/IP, Internet Protocol Exchange (IPX), or NetBEUI data packets and encrypts the data being transmitted as it is tunneled through the Internet. PPTP clients can connect to any Microsoft-compatible PPTP servers via the Internet with proper security credentials. This service, shipped with Windows XP Professional, allows users to connect to the Internet using local (nonlong-distance) connections and offers them a way to connect to PPTP computers in remote locations without incurring toll charges or requiring dedicated data lines.
- Layer 2 Tunneling Protocol (L2TP) An alternative to PPTP, L2TP was new to Windows 2000 and offers similar functionality to PPTP. However, L2TP is an industry-standard VPN protocol and is shipped with Windows XP Professional. L2TP also encapsulates TCP/IP, IPX, or NetBEUI data packets and encrypts the data being transmitted as it is tunneled through the Internet. You can also use L2TP in conjunction with Microsoft IP Security (IPSec) for enhanced security. L2TP is covered in more detail later in this chapter.

- IPSec This is a relatively new Internet security protocol, also referred to as *Secure IP*. It provides computer-level authentication in addition to data encryption for VPN connections that use the L2TP protocol. IPSec negotiates between the client computer and the remote tunnel server before an L2TP connection is established, which secures both authentication passwords and data. L2TP uses standard PPP-based authentication protocols, such as Extensible Authentication Protocol (EAP), Microsoft Challenge Handshake Authentication Protocol (MSCHAP), CHAP, Shiva Password Authentication Protocol (SPAP), and Password Authentication Protocol (PAP) with IPSec.
- World Wide Web (WWW) publishing service This is a major component of Internet Information Services (IIS), which ships with Windows XP Professional. Although not installed by default in Windows XP Professional, IIS and the WWW publishing service provide web page hosting for HTML-based and Active Server Pages (ASP)-based documents.
- File Transfer Protocol (FTP) service This is another major component of IIS. FTP is an industry-standard protocol for transferring files between computers over TCP/IP-based networks, such as the Internet.
- Simple Mail Transfer Protocol (SMTP) The Microsoft SMTP service implements the industry-standard SMTP to transport and deliver email messages. The SMTP service for Windows XP is also a component of IIS.

Understanding TCP/IP Computer Addresses

TCP/IP assigns a unique set of numbers to each computer that is connected to a TCP/IP-based network or internetwork. This set of numbers consists of four separate numbers, each delimited by a period or a dot (.). For example, an IP address of 192.168.1.20 illustrates this concept, known as dotted-decimal notation. Each device on a TCP/IP-based network must be assigned a unique IP address so that it can send and receive data with the other devices on the network. A network device can be a computer, a printer, a router, a firewall, and so on.

We write IP addresses in a dotted-decimal format for ease and convenience. However, TCP/IP addresses are actually 32-bit binary numbers! By converting these binary numbers into decimal, most of us can work with these addresses much more easily than if we had to work with them in their native binary format. The real binary address of 192.168.1.20, previously mentioned, translates into 11000000.10101000.00000001.00010100.

If you're not sure how to convert decimal numbers into binary or vice versa, just use the Windows Calculator by selecting Start, Run, typing **calc**, and clicking OK. Select View, Scientific and you can easily perform these conversions.

Certain IP addresses are reserved for specific functions:

- The address 255.255.255.255 (11111111.11111111.11111111.11111111 in binary) is reserved for *network broadcasts*.
- The IP address 127.0.0.1 (1111111.00000000.00000000.00000001 in binary) is reserved as a *loopback address* for testing proper configuration of the IP address(es) for the local host computer.
- The address schemes 192.168.x.y, 172.16.0.0 to 172.31.255.255, and 10.0.x.y have been *reserved as nonroutable* by the bodies that govern the Internet.

Therefore, IP addresses such as 192.168.1.20 and 10.0.0.7 are restricted to use only for the internal addressing of LANs. By definition, you cannot route these addressing schemes onto the Internet. Routers (devices that route network data packets) do not forward any data packets that originate with a nonroutable addressing scheme.

Understanding Classful IP Addressing

A look at Classful IP addressing takes us back to the beginning of TCP/IP itself. Classful addressing was adopted as RFC 791 and was the first major addressing scheme. Three address classes were used for typical network communication. These three ranges include A, B, and C class ranges. The difference between each class was the number of bits that made up the class prefix. For example, an IP address of 10.1.1.1 would be in the Class A range because the first octet, 10, starts with a prefix that is within the Class A range. Table 7.1 shows the different classes and the corresponding prefix ranges.

Table 7.1. Classful IP Address Prefix Ranges and Their Associated Address Classes

Prefix Range	**Address Class**
0127	A
128191	B
192223	C
224239	D
240255	E

You can quickly determine the class ranges by starting off with an octet of all zeros and turning on bits from the left-most part of the octet range. For a Class A range, the starting number would be an octet of all zeros up to an octet with the first bit turned on, 10000000. This would be a range of 0 to 127. The actual value of the octet with the first bit turned on is 128, which is the start of the next range. Then the end of the next range would be up to the second bit turned on, 11000000. So, the Class B would be a range of 128 to 191. If you follow this pattern for the remaining ranges, you will never be at a loss as to which range an IP address falls into.

When you see a reference to a Classful IP address scheme, it is referring to an address scheme that does not break up these classes. It is no surprise that because of the influx of the Internet, these address ranges are already purchased. If you want to get a range of IP addresses, you will need to contact an Internet service provider (ISP).

Understanding Variable-Length Subnet Masks

When you contact that ISP, you might be surprised that you can't obtain your own Class B range. What you might get instead is a portion of a Class B range.

When a Classful IP address range is broken down into smaller pieces, you need to use a *variable-length subnet mask* (VLSM), as shown in Figure 7.1. The standard subnet masks that come with the Classful IP ranges are as follows:

- Class A255.0.0.0
- Class B255.255.0.0
- Class C255.255.255.0

Figure 7.1. Running the ipconfig command with the /all option.

[View full size image]

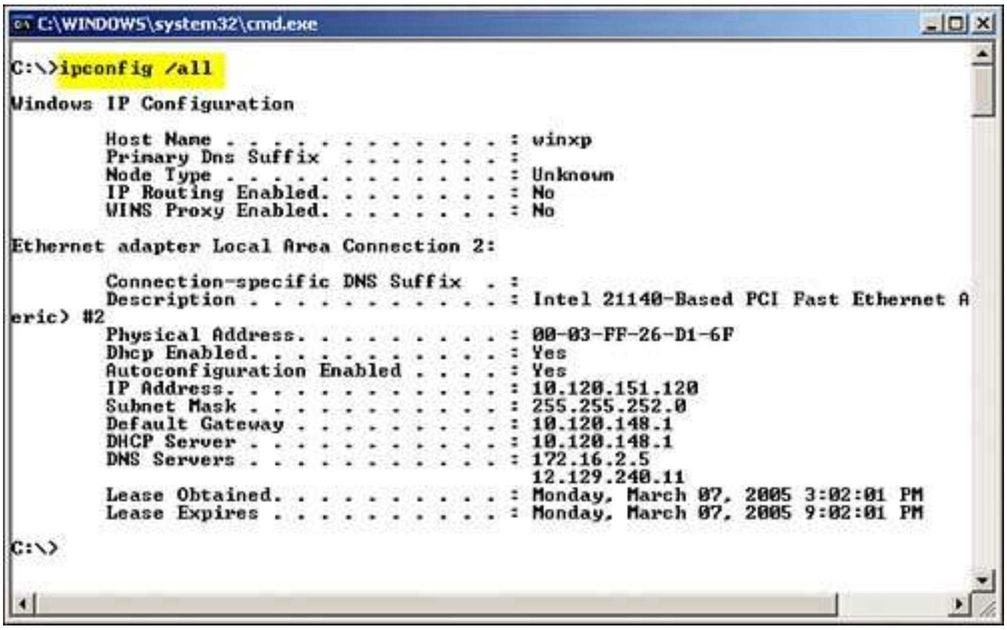

```
C:\>ipconfig /all

Windows IP Configuration

        Host Name . . . . . . . . . . . . : winxp
        Primary Dns Suffix  . . . . . . . :
        Node Type . . . . . . . . . . . . : Unknown
        IP Routing Enabled. . . . . . . . : No
        WINS Proxy Enabled. . . . . . . . : No

Ethernet adapter Local Area Connection 2:

        Connection-specific DNS Suffix  . :
        Description . . . . . . . . . . . : Intel 21140-Based PCI Fast Ethernet A
eric) #2
        Physical Address. . . . . . . . . : 00-03-FF-26-D1-6F
        Dhcp Enabled. . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . : Yes
        IP Address. . . . . . . . . . . . : 10.120.151.120
        Subnet Mask . . . . . . . . . . . : 255.255.252.0
        Default Gateway . . . . . . . . . : 10.120.148.1
        DHCP Server . . . . . . . . . . . : 10.120.148.1
        DNS Servers . . . . . . . . . . . : 172.16.2.5
                                            12.129.240.11
        Lease Obtained. . . . . . . . . . : Monday, March 07, 2005 3:02:01 PM
        Lease Expires . . . . . . . . . . : Monday, March 07, 2005 9:02:01 PM

C:\>
```

When you want to use only a portion of the address class range, you need to alter the standard subnet mask. A typical example of using VLSM is when you need to break up a Class C range into smaller ranges. A typical Class C range contains 254 IP addresses for hosts. Many smaller companies don't need this many addresses, so they will use a Class C address range that uses a VLSM to break up the range. This is done to create smaller pools of IP addresses. If you had a company that needed only 50 IP addresses, you could use a subnet mask of 255.255.255.192 with a Class C range of IP addresses. This subnet mask would break up the original Class C range and create four IP address ranges containing 62 IP addresses each.

Understanding Classless Interdomain Routing

The technology that we just looked at, VLSM, takes a Classful IP address range and makes more IP networks with fewer IP addresses. This is great for smaller companies or companies that want to break up the network into segments to reduce broadcasts. However, what if you are a larger company and you require additional IP addresses for one network segment? For this solution, you will need to combine IP address ranges. This is called Classless Interdomain Routing (CIDR). With CIDR, multiple subnets are seen as a single logical network of IP addresses. CIDR does have limitations, such as routing protocols and hardware devices. However, if your network can support CIDR, it just might be the solution that you are looking for.

When you are determining whether two IP address ranges can be combined with CIDR, you need to first determine if they share the higher-order bits. The following examples help explain how this works:

Example 1	**Bits**
10.1.2.0/24	00001010.00000001.00000010.00000000
10.1.3.0/24	00001010.00000001.00000011.00000000
Example 2	**Bits**
10.3.2.0/24	00001010.00000011.00000010.00000000
10.1.3.0/24	00001010.00000001.00000011.00000000

The first example can use CIDR because the first 23 bits are the same and therefore can be combined into a single network by the use of a classless network. This would be accomplished by using a new subnet mask of 255.255.254.0, which would result in a new subnet that would have a total of 510 host addresses. The second example will not work with a shortened CIDR subnet because only the first 14 bits are the same.

To get more information on CIDR, refer to http://www.petri.co.il/what's_cidr.htm on the Internet or review the Microsoft Windows Server 2003 Resource Kit, "Internet Protocol (IP) Addresses: Getting Started" section.

Configuring TCP/IP

TCP/IP is installed by default when you install Windows XP Professional. However, you can override this default setting if your network does not require it or if you will not be on a network with Active Directory. In addition, the protocol's default configuration is to *obtain an IP address automatically*. This means that the computer automatically requests a unique TCP/IP address for your network from a DHCP server. If no DHCP server is available, the operating system invokes APIPA to query the other computers that are currently powered on and connected to the network so that it can assign itself a unique IP address.

To work with TCP/IP, you need to become familiar with the following terms:

- Subnet mask This is essentially an IP address filter that gets applied to each unique IP address. The subnet mask determines which part of the IP address for a computer specifies the network segment where the computer is located, versus which part of the IP address specifies the unique host address for that individual computer. As an example, an IP address of 192.168.1.20 with a subnet mask of 255.255.255.0 is determined to have the network ID of 192.168.1. The host address for the computer, therefore, is 20. This is analogous to the street name of a postal address versus the actual house number of the address. The street may have many houses, but only one house has a house number of 20.
- Default gateway This IP address specifies the router for the local network segment (or subnet). If this address is absent, the computer cannot communicate with other computers that are located outside of the local network segment.

Default gateway information is often obtained through DHCP if the computer is configured to obtain an IP address automatically.

- Preferred and alternate DNS servers Having more than one DNS server on a network helps provide load balancing and fault tolerance for client computers that need to perform hostname-to-IP address lookups as well as IP address-to-hostname lookups. DNS is also used to find domain-based services such as domain controllers, DFS roots, and Global Catalog servers. Name resolution is a critical issue in TCP/IP. DNS server information is often obtained through DHCP if the computer is configured to obtain an IP address automatically.
- WINS server addresses WINS provides name resolution between NetBIOS computer names and IP addresses. WINS server addresses are often obtained through DHCP if the computer is configured to obtain an IP address automatically.

To manually set up a Windows XP Professional computer with a static IP address for the TCP/IP network protocol, click Start, Control Panel, Network Connections, and then select the Local Area Connection that you want to configure. Right-click the Local Area Connection icon and then click Properties. To configure the necessary settings so that TCP/IP can communicate with other computers and devices over the network, follow these steps:

1. Click Internet Protocol (TCP/IP) and then click Properties.

2. Click Use the Following IP Address.

3. Type the IP Address, Subnet Mask, and Default Gateway.

4. Type the proper IP address for a Preferred DNS Server and an Alternate DNS Server (if any).

5. Click the Advanced button to add additional IP addresses and default gateways. You can also add, edit, or remove DNS server address information, and you can change other DNS settings. You can specify IP addresses for any WINS servers on the network, enable NetBIOS name resolution using an LMHOSTS file, and enable or disable NetBIOS over TCP/IP. You can also set up IPSec and TCP/IP filtering as optional settings from the Advanced TCP/IP Settings properties sheet.

6. Click OK to close the Advanced TCP/IP Settings Properties dialog box.

7. Click OK to close the Internet Protocol (TCP/IP) Properties dialog box.

8. Click OK to close the Local Area Connection Properties dialog box.

Troubleshooting TCP/IP

Windows XP Professional comes with several software tools and utilities to help you isolate and resolve TCP/IP-related issues. You must run these utilities from the command line. Connectivity tools include the following:

- Finger Displays information about a user for a particular computer. The target computer must be running the Finger service.

- FTP Transfers files to and from FTP servers over a TCP/IP connection.
- LPR Sends one or more files to be printed via a line printer daemon (LPD) printer.
- RCP Copies files between a Windows XP Professional computer (or a computer running Windows Server 2003) and a computer system running the remote shell daemon (RSHD). Windows 2000 Professional and Windows XP clients cannot run the RSHD daemon, but Unix systems can.
- REXEC Executes commands on remote computer systems that are running the REXEC service. Windows 2000, Windows XP, and Windows Server 2003 systems all have the capability to run the REXEC service.
- RSH Executes commands on remote computer systems that are running the RSH service. Windows 2000, Windows XP, and Window Server 2003 systems do not run the RSH service.

TIP

A utility included with the Windows 2000 Server Resource Kit enables the RSH service to run on a Windows 2000 system. The utility is called **RSHSVC.EXE**. The Windows Server 2003 Resource Kit does not include this utility.

-
- Telnet Establishes a terminal emulation session for working on remote systems, including environments such as Unix, Mainframe, and minicomputers.
- Trivial File Transfer Protocol (TFTP) Copies files to and from remote computers that are running the TFTP service.

Diagnostic tools include the following:

- Address Resolution Protocol (ARP) Lists and edits the IP-to-Ethernet (or Token Ring) physical translation tables that ARP uses.
- HOSTNAME Lists the name of the local host (computer).
- IPCONFIG Shows all current TCP/IP configuration settings for the local computer, such as its IP address, subnet mask, and any WINS servers and DNS servers assigned to the computer.

ALERT

Special optional parameters for the **IPCONFIG** command deal with the DNS portion of the IP session. These include the **/registerdns, /displaydns**, and **/flushdns** switches. These switches register the client with DNS, show the current DNS cache, and flush out the DNS cache, respectively. You can use the **IPCONFIG** command to release and renew an automatically assigned IP address from a DHCP server with the **/release** and the **/renew** switches.

-
- LPQ Shows the current status of the print queue on a computer that is running the LPD service.
- NBTSTAT Delineates network protocol statistics and lists the current connections that are using NetBIOS over TCP/IP.

ALERT

Don't forget about the **R** and **RR** switches that can help refresh the cache, send release packets to WINS, and then refresh the client connection.

-

- NETSTAT Delineates network protocol statistics and lists the current TCP/IP connections.
- PING Used to test TCP/IP-related connectivity to remote computers. This command also verifies the proper TCP/IP configuration of the local host computer by attempting to ping the loopback address for the local host (computer). For example: `ping 127.0.0.1`
- ROUTE Edits the local computer's routing tables.
- TRACERT Displays the route (path) that data packets follow as they travel from the local computer to a remote destination computer.

Troubleshooting TCP/IP Configuration and Connectivity

Whenever you initially set up TCP/IP, you should always test and verify that the protocol is working properly. Following are the steps you can take to check the computer's TCP/IP configuration and test its connectivity:

1. Open a command prompt window (`ipconfig` and `ping` are command-line-only utilities).

2. Run `ipconfig` to display the computer's current IP configuration. Use `ipconfig /all` to display more detailed information, as shown in Figure 7.1.

3. Use the `ping` command to ping a computer's loopback address, which is 127.0.0.1, by default (for example, `ping 127.0.0.1`). This tests whether TCP/IP is correctly installed and properly bound to the network adapter card(s). Figure 7.2 shows the response from pinging the loopback IP address.

Figure 7.2. Running the* ping *command using the computer's loopback IP address.

[View full size image]

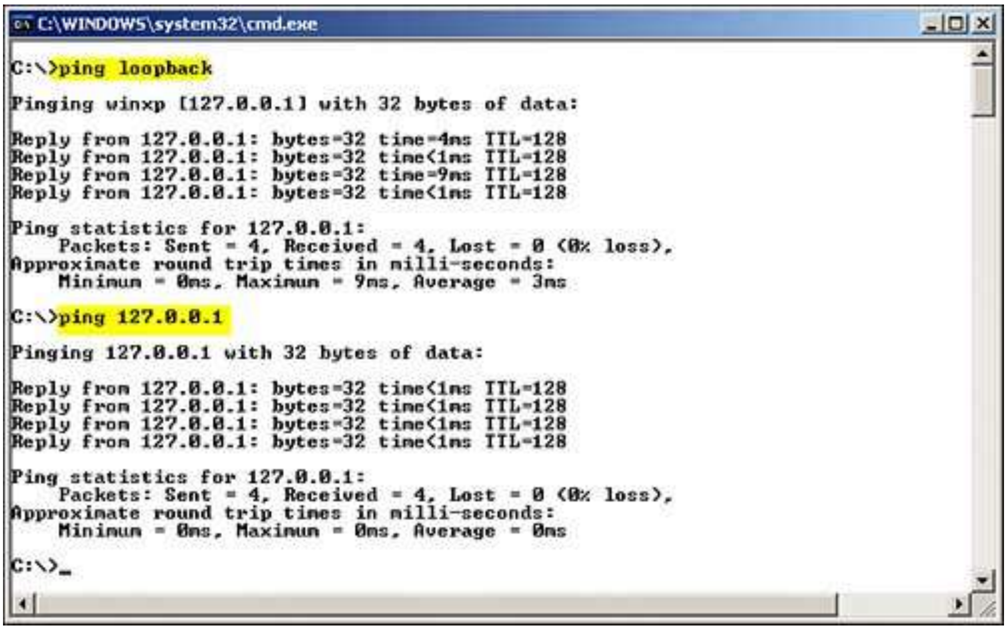

```
C:\>ping loopback

Pinging winxp [127.0.0.1] with 32 bytes of data:

Reply from 127.0.0.1: bytes=32 time=4ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time=9ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128

Ping statistics for 127.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 9ms, Average = 3ms

C:\>ping 127.0.0.1

Pinging 127.0.0.1 with 32 bytes of data:

Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128

Ping statistics for 127.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\>_
```

4. Ping the IP address of the local computer to verify the uniqueness of the IP address on the network.

5. Ping the IP address of the default gateway for the local subnet to check that the default gateway is up and running. This step also demonstrates whether the computer can successfully communicate over the local network segment.

6. Ping the IP address of a computer that is located on a different network segment. This step indicates whether the computer can send and receive network data packets through a router.

Using APIPA

If a computer is set up to obtain an IP address automatically from a DHCP server, but no DHCP servers are available, APIPA temporarily assigns an IP address to the local computer while it searches the network to make sure that no other network devices have been assigned the same IP address. By running `ipconfig`, you can view the current TCP/IP information for the local computer. An address such as 169.254.x.y generally indicates that APIPA is currently in effect.

Windows XP clients do not indicate that they are unable to obtain an IP address from a DHCP server. By default, instead of a notification, a Windows XP computer obtains an APIPA address without any warning or message. It is essential that the help desk and IT staff be made aware that this is the default behavior and that they can be sure to add it to their list of items for troubleshooting TCP/IP communication issues.

Configuring WINS replication

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Configuring WINS replication

Before configuring replication, you should first carefully design and review your WINS replication topology. For wide area networks (WANs), this planning can be critical to the success of your deployment and use of WINS. In general, WINS offers you the following options and possibilities to choose from when you are configuring replication:

- You can manually configure WINS replication for a WAN environment.
- For some larger campus networks, you might also configure WINS to replicate within a LAN environment.
- In smaller or bounded LAN installations, you might consider enabling and using WINS automatic partner configuration for simplified setup of WINS replication.

- In some larger or more global installations, you might have to configure WINS across untrusted Windows NT domains.

Replication across wide area networks

When configuring WINS replication, the two most important planning issues are:

- Whether your WINS replication occurs over slower WAN links.
- The length of time required for all replicated changes in the WINS database to converge and achieve consistency on the network.

The frequency of WINS database replication between WINS servers is a major planning issue. The WINS server database should be replicated frequently enough to prevent the downtime of a single WINS server from affecting the reliability of the mapping information in other WINS servers. However, the time interval between replications should not be so small that it interferes with network throughput.

Network topology can influence your decision on replication frequency. For example, if your network has multiple hubs connected by relatively slow WAN links, you can configure WINS database replication between WINS servers on the slow links to occur less frequently than replication on the local area network or on fast WAN links. This reduces traffic across the slow link and reduces contention between replication traffic and WINS client name queries.

For example, as shown in the following figure, the intervals used for configuring pull replication for WINS servers could be set according to proximity between servers and the likelihood of a high-speed link connecting servers configured as replication partners of each other.

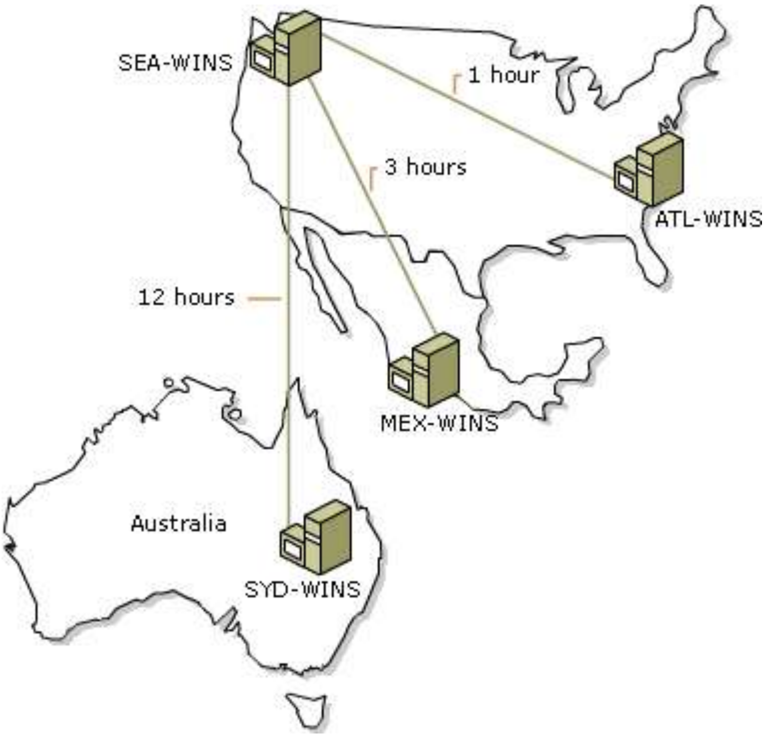

In this example, two WINS servers--SEA-WINS and ATL-WINS--are located in Seattle and Atlanta. Another two servers--MEX-WINS and SYD-WINS--are located outside the United States in Mexico City and Sydney. Because SEA-WINS and ATL-WINS are located across a higher-speed WAN link, they are configured to replicate once every hour. For the servers outside the United States, replication intervals might be set higher, such as either 3 or 12 hours, because lower-speed links are associated with these connections.

Planning the convergence time

Convergence time is the time needed to replicate a new entry in a WINS database from the WINS server that owns the entry to all other WINS servers on the network. When planning for WINS servers, you must decide what is acceptable as the convergence time for your network.

Using the previous example, if a WINS client in Seattle registers its name with the WINS server in the figure labeled SEA-WINS, other WINS clients in remote locations might query for the client name and not find it in WINS. In some cases, if the address of the client has been changed, the remote locations would only have a mapping for the previous IP address of the client in their local server databases.

WINS clients that query any of the other WINS servers in the figure (that is, ATL-WINS, MEX-WINS, SYD-WINS) only receive a positive and correct response when the new or updated mapping for the client has been fully replicated from the owning WINS server, SEA-WINS, to all of the other servers. Replications from SEA-WINS could be triggered in one of two ways:

- A push replication trigger is sent based on the frequency of updates made at SEA-WINS.
- A pull replication trigger is sent to other WINS servers based on the configured **Replication interval** set in **Replication Partners Properties** at SEA-WINS.

Given this configuration, a new or updated entry at SEA-WINS is only replicated when the pull replication interval expires, but queries at remote sites for the new name mapping might not succeed until later. This is because the time interval for replication to ATL-WINS is one hour; it is three hours for MEX-WINS; and it is 12 hours for SYD-WINS. In this case, WINS convergence time is calculated as:

```
12 hours + 2 * (1 hour) = 14 hours
```

In reality, name query requests can succeed before the convergence time elapses. This happens when the entries replicate over a shorter path than the worst-case path. It also happens when the **Number of changes in version ID before replication** threshold in push replication settings passes before the **Replication Interval** expires in pull replication settings. This results in earlier replication of the new entry. The longer the replication path, the longer the convergence time.

Notes

- After each administrator configures the respective WINS server, you can use **Replicate Now** (in the WINS console) to start immediate replication between the configured WINS servers.
- When replication occurs, you can use **Display Records** to view records for the remote WINS server database. However, to make changes to the remote WINS database, you must be logged on under an account with required user rights in the respective domains for both servers involved.
- Replication can also occur between two separate stand-alone servers that participate in a workgroup.

Replication across local area networks

When configuring WINS replication for local area networks (LANs), the issues are similar to those that occur in WAN environments, although less critical.

Because the speed of the underlying network links for LANs are much greater than for WANs, it might be acceptable to increase the frequency of WINS database replication, by optimizing push and pull parameters for LAN-based replication partners. For push/pull partners, you can do this by decreasing the **Number of changes in version ID before replication** and **Replication interval** settings accordingly, from what you would use for WAN-based partners on slower links.

For example, between LAN-based replication partners it often works to enable WINS to use a persistent connection between the servers. If a persistent connection was not used, the normal update count threshold would default to a minimum of 20. With persistence, a smaller update count value can be specified.

Next, you could specify a much smaller number, such as a value of one to three in **Number of changes in version ID before replication** before WINS sends a push replication trigger to the other partner. For pull partners, you might also consider setting the **Replication interval** to a value in minutes, instead of hours.

As in WAN replication planning, the WINS server database should be replicated frequently enough to prevent the downtime of a single WINS server from affecting the reliability of the mapping information in other WINS servers. However, the time interval between replications should not be so small that it interferes with network throughput.

For network-intensive environments, it is a good idea to use a network monitoring tool such as Network Monitor to help measure and determine how to optimize your WINS replication strategy.

Automatic partner configuration

You can configure a WINS server to automatically configure other WINS server computers as its replication partners. With this automatic partner configuration, other WINS servers are discovered when they join the network and are added as replication partners.

Automatic configuration is possible because each WINS server announces its presence on the network through periodic multicasts. These announcements are sent as IGMP messages for the multicast group address of 224.0.1.24 (the well-known multicast IP address reserved for WINS server use).

Automatic replication configuration monitors multicast announcements from other WINS servers and automatically performs the following configuration steps:

- Adds the IP addresses for the discovered servers to its list of replication partner servers.
- Configures the discovered servers as both push and pull partners.
- Configures pull replication at two-hour intervals with the discovered servers.

If a remote server is discovered and added as a partner through multicasting, it is removed as a replication partner when WINS shuts down properly. To have automatic partner information persist when WINS restarts, you must manually configure the partners.

To manually configure replication with other WINS servers, use the WINS console to specify roles for each partner and any related information.

Notes

- Automatic partner configuration is typically useful in small networks, such as single subnet LAN environments. It can, however, be used in routed networks. For WINS multicast support in routed networks, the forwarding of multicast traffic is made possible by configuring routers for each subnet to forward multicast traffic for the reserved WINS server multicast group. The destination IP address for this group is 224.0.1.24 between routed subnets.
- Because periodic multicast announcements between WINS servers can add traffic to your network, automatic partner configuration is recommended only if you have a small number of installed WINS servers (typically, three or fewer) on the reachable network.

Replication between untrusted domains

It is possible to set up WINS replication between one or more WINS servers in domains that do not have a trust relationship. You can do this without a valid user account in the untrusting domain. To configure replication, an administrator for each WINS server must use the WINS console to manually configure that server to permit this replication.

WINS (Windows Internet Name Service)

Home > Protocols > WINS   Update: 2005-11-25 15:14:27   I have words to say about this protocol

On this page

SUMMARY

Protocol	:	Windows Internet Name Service
Protocol suite	:	TCP/IP
Layer	:	Application Layer
Ports	:	1512 (TCP, UDP)

DESCRIPTION

When using TCP/IP to communicate on a network, the friendly computer name that is used in a net user command, must be resolved to an IP address. This is necessary because TCP/IP does not know how to establish communication with a computer name, such as \server1, but does know how to communicate with 223.223.223.1. In order to resolve the computer name to its IP address, TCP/IP can use a variety of methods: broadcasts, a static mapping file (LMHOSTS), or a name server (WINS). The Windows Internet Name Service (WINS) was designed to eliminate the need for broadcasts to resolve computer names to IP addresses and provide a dynamic database that maintains computer name to IP address mappings. There are two types of systems that use WINS:

WINS Clients

WINS Clients are configured with the IP address of one or more WINS Servers. On startup, WINS Clients communicate directly with a WINS Server to register their computer name and corresponding IP address. When a WINS Client needs to resolve a computer name to an IP address, such as when a net use \servershare is performed, the WINS Client sends a request to the WINS Server for the IP address for the computer name being used.

WINS Servers

A WINS (Windows Internet Name Service) Server maintains a database that maps the IP addresses of WINS Clients to their computer name, also referred to as a NetBIOS name. Therefore, instead of using broadcasts to resolve a computer name to an IP address when trying to establish a network connection, WINS Clients request the IP address for the desired system from a WINS Server which retrieves the IP address from its database.

WINS has been enhanced for the release of Microsoft Windows 2000 Server. The result is an easier-to-manage and more robust solution for mapping NetBIOS names to IP addresses on Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

Windows 2000 WINS includes server enhancements, additional client functions, and an improved management tool. WINS provides a distributed database for registering and querying dynamic computer name-to-IP address mapping in a routed network environment. This support for dynamic registering of NetBIOS computer names means that WINS can be used with Dynamic Host Configuration Protocol (DHCP) services to provide easy configuration and administration of Windows-based TCP/IP networks.

The WINS server solves the problems inherent in resolving NetBIOS names through IP broadcasts, and frees network administrators from the demands of updating static mapping files, such as LMHOST files. WINS, which is compliant with the NetBIOS Name Server (NBNS) RFC s (1001/1002), also automatically updates the WINS database when dynamic addressing through DHCP results in new IP addresses for computers that move between subnets. Neither the user nor the network administrator needs to make manual accommodations for such name resolutions.

The new implementation of WINS provides a number of features, including:

- **Persistent connections**
  This configurable feature allows each WINS server to maintain a persistent connection with one or more replication partners to eliminate the overhead of opening and terminating connections and to increase the speed of replication.

- **Manual tombstoning**
  Use of the Manual tombstoning feature marks a record for deletion so that the tombstone state for the record is replicated across all WINS servers, preventing an undeleted copy of the

record on a different server database from being re-propagated.

- **Improved management tools**
  The WINS Manager is fully integrated with the Microsoft Management Console (MMC), providing a more user-friendly and powerful environment for viewing and managing WINS information.

- **Enhanced filtering and record searching**
  These functions help locate records of interest by showing only those that fit a specific criteria. This is particularly useful for analyzing very large WINS databases.

- **Dynamic record deletion and multi-select**
  Managing the WINS database is made easier with dynamic record deletion and multi-select. Dynamic and static records can be deleted, and the point-and-click interface makes it possible to delete files with non-alphanumeric characters that could not be handled from the command line.

- **Record verification and version number validation**
  Two tools are available for quickly checking the consistency between various WINS servers. The tests are done by comparing the IP addresses of a NetBIOS name query returned from different WINS servers or by examining owner address to version-number mapping tables.

- **Export function**
  The Export command can be used to place WINS data into a comma-delimited text file that can be imported into Microsoft Excel, reporting tools, scripting applications, and so on, for analysis and reporting.

- **Increased fault tolerance**
  Windows 2000 and Windows 98 allow a client to specify more than two WINS servers (up to a maximum of 12 addresses) per interface. The extra WINS server addresses are used only if the primary and secondary WINS servers fail to respond.

- **Dynamic re-registration**
  WINS clients can now re-register their NetBIOS name-to- IP address mapping without rebooting the server.

All of this combines to make Windows 2000 WINS a superior choice for NetBIOS name resolution. The new generation of WINS services are designed to make many network management tasks much easier for network managers.

How does WINS work?

As soon as a WINS client obtains the IP configuration from the DHCP server (or on bootup if the WINS server's IP address was statically assigned), the WINS client issues a NameRegistrationRequest message to the WINS server. Unlike standard NetBIOS behavior, this message isn't a broadcast. It is a message sent only to the primary WINS server and includes the clienti¯s computer name and IP address.

The WINS server checks to see whether the computer name is listed in its database. If it isn't listed, the WINS server assumes that it's unique on the network and responds with a positive WINS name registration response. The registration response includes a period called the time-to-live (TTL) during which the registration is valid. If the name isn't unique, a negative response is sent to the client, and the WINS server sends a challenge to the name's current owner. Typically, the computer that currently owns the name acknowledges that it is alive on the network and the negative response message informs the new computer that there is a conflict.

To configure DHCP and WINS in an NT network, DHCP options 44 and 46 must be configured. Recall from Chapter 8 that option 44 is the IP address of the WINS/NBNS server, and option 46 is the Node Type of the WINS/NBNS server. As we explained in Chapter 6, there are four possible node types, three of which are applicable in this instance:

P node (point-to-point)

M node (mixed - broadcasts then point-to-point)

H node (hybrid - point-to-point then broadcast)

H node is the default configuration and is suitable for most networks. If you have a small network or if your WINS server is located across a router, you may want to consider M node. This may provide some network-traffic optimization if most of your resources are local and not across the router.

Advantages and disadvantages

There are many advantages to using WINS. It provides dynamic NetBIOS name resolution, thereby reducing or eliminating the administrative effort required to update HOSTS, LMHOSTS, and/or DNS files. WINS reduces the number of IP broadcasts, thus reducing network traffic. It also provides a centralized management scheme for a NetBIOS computer name database, and it allows replication of that database to other WINS servers.

The disadvantages of using WINS include

- WINS is a proprietary Microsoft service. Only Windows-based computers or computers that understand SMB (Server Message Block) networking or Common Internet File System (CIFS) can be clients. Your options include computers running NT Server, NT Workstation, Windows 95/98, Windows for Workgroups, LAN Manager 2.x, or systems running products like SAMBA. However, a WINS proxy can be used to help resolve NetBIOS name queries for non-WINS enabled computers.

- Usually, only NT servers can be WINS servers. However, you can create a WINS server using SAMBA; this method is discussed in "WINS Servers - Unix and Linux" later in this chapter.

- WINS adds one more level of complication to the Microsoft name resolution service. 2000 Server is striving to eliminate the need for WINS by adopting dynamic DNS and encouraging the elimination of NetBIOS traffic. However, because many applications use NetBIOS, eliminating NetBIOS is probably not a practical solution in the short term.

Setting up WINS servers

The following sections describe WINS server setup and operation in NT 4.0 networks, 2000 Server networks, and heterogeneous Unix/NT networks.

- **WINS Servers - NT 4.0**
  Normally only NT 4.0 and 2000 Server computers can be set up as WINS servers. To add the WINS service to an NT 4.0 computer, click Start, and select Settings, Control Panel, and Network. Switch to the Services tab and click Add. Select Windows Internet Naming Service, insert the NT CD-ROM and select {drive letter}/I386. (If you are installing on an Alpha-processor-based server, go to the Alpha directory instead of I386.) Only administrators, by default, can add the WINS (or any) service to

an NT computer. After you install WINS, reboot the computer and you are in business. Once the WINS service is installed, your Administrative Tools list contains WINS Manager option, and you can use this to configure WINS.

- **WINS Servers - 2000 Server**
  WINS servers are typically not required with 2000 Server because most of the NetBIOS-related services are gone. However, if you have clients like Windows 95, Windows for Workgroups, NT Workstation and Server 3.x and 4.0, you may want to set up a WINS server on your 2000 Server. Click Start and select Settings, Control Panel, Add/Remove Programs, Configure Windows. Click Components and select Networking Options. Check the WINS server box.

  With NT 4.0, when the statistics pane of WINS Manager window is grayed, the WINS server isn't running. If the service is running, it shows the server start time. WINS Manager user interface in 2000 Server uses the Microsoft Management Console (MMC) and, as a result, is more intuitive. For example, it is easier to determine the server status.

- **WINS Servers - Unix and Linux**
  One of the most popular and well-written networking programs is SAMBA, an SMB server for Unix and Unix-like operating systems. Versions of SAMBA run on VMS, NetWare, MVS, MPE/ix, and other operating systems. SAMBA can be configured as a WINS server. Typically this is recommended only if you don't use an NT Server as a primary domain controller (PDC). In situations where you have a large number of Windows computers but no NT Servers - for example, a remote office or a department that won't use NT Servers - you could consider using this free product.

  After you install SAMBA, all you have to do to have a Unix-based WINS server is edit the smb.conf file, usually located in the /etc directory. Look for the entry "wins support = yes." By default this is a comment. Uncomment this line by removing the semicolon, and you have a WINS server that runs on a Unix computer. Make sure that the entry "wins server = a.b.c.d" remains a comment because SAMBA can't be a WINS server and a WINS client simultaneously.

Setting up WINS clients
If you are configuring WINS clients using a DHCP server, you don't have to worry about the following steps. Configuring clients using DHCP also avoids the requests from NT for a reboot whenever network parameters are changed. Just set up the DHCP server options 44 and 46 via the DHCP Manager menu item DHCP Options. Select the DHCP server, DHCP Options, and the appropriate scope. Choose Scope to configure options for one scope only; choose Global to configure options for all scopes. Select the options and the values in the resulting screen.

To force the client to surrender its current IP address and renew the newly configured IP address, go to the client and use the ipconfig/release and ipconfig /renew commands at the command prompt (or from the graphical utility winipcfg.exe in the case of Windows 95 or 98 computers).

- **WINS Clients - NT 4.0**
  To use WINS, NT 4.0 clients must be configured with the address of a WINS server. If you are using DHCP, the clients are configured automatically, but you can also assign an address. To configure an NT 4.0 client to use WINS, open the Control Panel, select Network, Protocols, TCP/IP, and switch to the WINS Address tab (Figure 9.1). Specify the primary and secondary WINS servers and, if you wish, enable DNS lookup or LMHOSTS lookup. If DNS lookup is selected, DNS will be used if WINS fails to resolve the name. The resolution is performed by joining the domain name with the computer name to form a fully-qualified domain name. If LMHOSTS is selected, LMHOSTS will be used to resolve a query if the WINS query fails. If both are selected, the order of name resolution attempts for the host name will be as follows:

  - Check local computer host name.
  - Check the HOSTS file.
  - If there is no entry in the HOSTS file and if a DNS server is available, query the DNS server.
  - If the DNS server doesn't respond and the resource was accessed using a name, try NetBIOS name resolution.
  - If the resource is being accessed through an IP address, send a NetBIOS request (NetBIOS Adapter Status Request) to the IP address for a list of NetBIOS names registered for that adapter.

- Parse the results for a computer name.

On the other hand, if a NetBIOS name resolution is being attempted, the order tends to be as follows:

- Check the NetBIOS local name cache.
- Check the NetBIOS name server (WINS).
- Send broadcasts within an IP subnetwork.
- Check the LMHOSTS file(s).
- Check the HOSTS file.
- Query the DNS server.

- **WINS Clients - 2000 Server**
  If you have a mixed NT 4.0 and 2000 Server network or if you are using WINS/NBNS servers (for example, SAMBA) on other platforms, you may need to configure 2000 Server clients (workstations and servers) to be registered and resolved through WINS servers. Open the Control Panel, select Network Connections, and right-click Local Area Connections. From the menu, select Properties, Internet Protocol, Advanced, and switch to the WINS tab. Add the IP address of your WINS server and you can use the specified WINS server for NetBIOS name resolution.

- **WINS Clients - Unix and Unix-like Systems**
  It's very easy to configure any computer running SAMBA as a WINS client, but recall from the server discussion that SAMBA can't be a WINS server and a WINS client at the same time. So, first ensure that the smb.conf file entry "wins support = yes" (which configures the SAMBA computer as a WINS server) is a comment (the default). Then edit the next line to read "wins server = www.xxx.yyy.zzz ," where www.xxx.yyy.zzz is the IP address of your WINS server.

  You don't have to reboot the Unix computer. SAMBA automatically reads the configuration file changes. To force the changes to take place immediately, rather than waiting for SAMBA to read the changes from the configuration file, you can stop and restart the SAMBA programs using the /etc/rc.d/init.d/smb stop and

/etc/rc.d/init.d/smb start commands.

Top of Page

EXAMPLES

Top of Page

PROTOCOL RELATIONS

- 🟥 Parent layer
- 🟩 Child layer

TCP/UDP WINS

Top of Page

GLOSSARY

Broadcast
Broadcast is the term used to describe communication where a piece of information is sent from one point to all other points. Broadcasting is a useful feature in e-mail systems. It is also supported by some fax systems.

In networking, a distinction is made between broadcasting and multicasting. Broadcasting sends a message to everyone on the network whereas multicasting sends a message to a select list of recipients.

Client
Clinet is a program which requests services of another program. It is a client part of a client-server architecture. Typically, a client is an application that runs on a personal computer or workstation and relies on a server to perform some operations. For example, an e-mail client is an application that enables you to send and receive e-mail.

Command
Command is an instruction to a computer or device to perform a specific task. Commands come in different forms. They can be: special words (keywords) that a program understands, function keys
choices in a menu and buttons or other graphical objects on your screen

Every program that interacts with people responds to a specific set of commands. The set of commands and the syntax for entering them is called the user interface and varies from one program to another.

DNS

DNS(Domain Name System or Service or Server), an Internet service that translates domain names into IP addresses. Because domain names are alphabetic, they're easier to remember. The Internet however, is really based on IP addresses. Every time you use a domain name, therefore, a DNS service must translate the name into the corresponding IP address. For example, the domain name www.example.com might translate to 198.105.232.4.

The DNS system is, in fact, its own network. If one DNS server doesn't know how to translate a particular domain name, it asks another one, and so on, until the correct IP address is returned.

Database

A database is an organized collection of data. The term originated within the computer industry, but its meaning has been broadened by popular use, to the extent that the European Database Directive (which creates intellectual property rights for databases) includes non-electronic databases within its definition. This article is confined to a more technical use of the term; though even amongst computing professionals, some attach a much wider meaning to the word than others.

Dynamic

Dynamic refers to actions that take place at the moment they are needed rather than in advance. For example, many programs perform dynamic memory allocation, which means that they do not reserve memory ahead of time, but seize sections of memory when needed. In general, such programs require less memory, although they may run a little more slowly.

Host

Host is a computer system that is accessed by a user working at a remote location. Typically, the term is used when there are two computer systems connected by modems and telephone lines. The system that contains the data is called the host, while the computer at which the user sits is called the remote terminal.

Host can refer to a computer that is connected to a TCP/IP network, including the Internet. Each host has a unique IP address.

Host can refer to provide the infrastructure for a computer service too. For example, there are many companies that host Web servers. This means that they provide the hardware, software, and communications lines required by the server, but the content on the server may be controlled by someone else.

IP address
IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IP address is a 32-bit numeric address written as four numbers separated by periods. Each number can be zero to 255. For example, 1.160.10.240 could be an IP address. Within an isolated network, you can assign IP addresses at random as long as each one is unique. However, connecting a private network to the Internet requires using registered IP addresses (called Internet addresses) to avoid duplicates.

The four numbers in an IP address are used in different ways to identify a particular network and a host on that network. Four regional Internet registries -- ARIN, RIPE NCC, LACNIC and APNIC -- assign Internet addresses from the following three classes.
Class A - supports 16 million hosts on each of 126 networks
Class B - supports 65,000 hosts on each of 16,000 networks
Class C - supports 254 hosts on each of 2 million networks

The number of unassigned Internet addresses is running out, so a new classless scheme called CIDR is gradually replacing the system based on classes A, B, and C and is tied to adoption of IPv6.

Internet
A global network connecting millions of computers. More than 100 countries are linked into exchanges of data, news and opinions.

Unlike online services, which are centrally controlled, the Internet is decentralized by design. Each Internet computer, called a host, is independent. Its operators can choose which Internet services to use and which local services to make available to the global Internet community. Remarkably, this anarchy by design works exceedingly well.

There are a variety of ways to access the Internet. Most online services, such as America Online, offer access to some Internet services. It is also possible to gain access through a commercial Internet Service Provider (ISP).

Linux
Linux is a freely-distributable open source operating system that runs on a number of hardware platforms. The Linux kernel was developed mainly by Linus Torvalds. Because it's free, and because it runs on many platforms, including PCs and Macintoshes, Linux has become an extremely popular alternative

to proprietary operating systems.

MMC
MMC (Microsoft Management Console) is an extensible common presentation service for management applications. MMC is included in the Windows&reg; 2000 operating system.

Microsoft
Microsoft founded in 1975 by Paul Allen and Bill Gates, Microsoft Corporation is one of the largest and most influential companies in the personal computer industry. In addition to developing the de facto standard operating systems -- DOS and Windows -- Microsoft has a strong presence in almost every area of computer software, from programming tools to end-user applications.

Network
Network is a group of two or more computer systems linked together. There are many types of computer networks, including: LANs (local-area networks), WANs (wide-area networks), CANs (campus-area networks), MANs (metropolitan-area networks) and HANs (home-area networks).

In addition to these types, the following characteristics are also used to categorize different types of networks: Topology, protocol and architecture.

SMB
Server Message Block (SMB) is a message format used by DOS and Windows to share files, directories and devices. NetBIOS is based on the SMB format, and many network products use SMB. These SMB-based networks include LAN Manager, Windows for Workgroups, Windows NT, and LAN Server. There are also a number of products that use SMB to enable file sharing among different operating system platforms. A product called Samba, for example, enables UNIX and Windows machines to share directories and files.

Server
A computer or device on a network that manages network resources. For example, a file server is a computer and storage device dedicated to storing files. Any user on the network can store files on the server. A database server is a computer system that processes database queries. Servers are often dedicated, meaning that they perform no other tasks besides their server tasks. On multiprocessing operating systems, however, a single computer can execute several programs at once. A server in this case could refer to the program that is managing resources rather than the entire computer.

TCP/IP

TCP/IP(transmission Control Protocol/Internet Protocol) is the suite of communications protocols used to connect hosts on the Internet. TCP/IP uses several protocols, the two main ones being TCP and IP. TCP/IP is built into the UNIX operating system and is used by the Internet, making it the de facto standard for transmitting data over networks. Even network operating systems that have their own protocols, such as Netware, also support TCP/IP.

TTL

TTL (Time to Live) is a field in the Internet Protocol (IP) that specifies how many more hops a packet can travel before being discarded or returned.

Unix

Unix or UNIX is a computer operating system originally developed in the 1960s and 1970s by a group of AT&T Bell Labs employees including Ken Thompson, Dennis Ritchie, and Douglas McIlroy. Today's Unix systems are split into various branches, developed over time by AT&T, several other commercial vendors, as well as several non-profit organizations.

WINS

WINS (Windows Internet Naming Service) is Microsoft's implementation of NetBIOS Name Server (NBNS) on Windows, a name server and service for NetBIOS computer names. Effectively, it is to NetBIOS names what DNS is to domain names - a central store for information, so that when a client needs to contact a computer on the network, it can query the central database for the address to contact rather than broadcasting a request for which address to contact.

20.3. NetBIOS for TCP/IP Name Service and Windows Internet Name Service

NetBT name service is the NetBT service used to translate NetBIOS names to IP addresses. There are two ways to get name service under NetBT: a broadcast mechanism in which each machine keeps a database, and a unicast mechanism in which there is a designated server. In theory, that designated server is called a NetBT Name Server, or NBNS. However, Microsoft's NBNS implementation is called Windows Internet Name Service (WINS). It's rare to see the general term "NBNS" used outside of standards documentation, even for non-Microsoft servers, which are technically implementations of NBNS, not WINS.

In order to minimize the inevitable confusion, we will call broadcast-based NetBT name service "NetBT name service", and unicast-based service "WINS". This reflects common usage and is no more arbitrary and confusing than any other naming scheme.

A NetBIOS name is up to 15 characters long.[121] NetBIOS names are "unqualified" (according to Microsoft documentation) which means two things:

[121]Technically speaking, all NetBIOS names are exactly 16 characters long, and Microsoft reserves the 16th character for administrative use. Short names are automatically and transparently padded to 16 characters. To the user, this is indistinguishable from having names up to 15 characters long.

- They must be unique within a server's area of control.
- They cannot contain periods.

There is no notion of hierarchy in NetBIOS. Only one machine can be named "foo", and there is no way to indicate in a name what server is responsible for it. In order to reduce problems with name collisions, Microsoft provides the concept of a NetBIOS *scope*, which strongly resembles a NIS domain; this is a string that defines a group of machines that talk to each other. If you are using NetBIOS scopes, a machine's name needs to be unique only within the scope. By default, machines use a null NetBIOS scope. The NetBIOS name and scope taken together can be up to 255 characters long (effectively limiting the scope to 240 characters), and the NetBIOS scope can contain periods.

NetBIOS scopes are much more limiting than NIS domains. Machines that are in different NetBIOS scopes cannot speak any NetBIOS protocols to each other, including file and printer sharing. On controlled networks, this can actually be an advantage, as it provides a small security improvement; the NetBIOS scope setting effectively acts as a password for network access. This is a protection from accidental misconfiguration, not from hostile action. The NetBIOS scope is passed in cleartext across the network as part of the NetBIOS hostname, and any attacker can simply read it from valid packets (it is particularly simple because it is sent in broadcast packets, so no special ability to snoop the network is required.)

It is important to keep in mind that WINS and NetBT name service are merely variants on the same service. Clients that use broadcast resolution run their own name servers, and although they expect only broadcast-based queries, they will respond to unicast queries. Do not assume that you are safe from remote requests just because you have avoided running WINS; normal NetBIOS clients are still running name servers that will feed their own data and any other data that they have cached to anybody who asks them. As we will see later, even machines that use WINS exclusively must have servers running that will respond to unicast name queries in order to have full WINS functionality. This is one of the things that the "Server" service normally does on clients.

Even though these services are tightly interrelated, the same machine may run both servers, in which case the WINS server will get unicast packets and the NetBT name server will get broadcast packets (Unix aficionados will probably find this upsetting). This means that the two services will run completely independently and may have different data on them. A WINS server will not provide data gathered by WINS to a NetBT name service client, or data gathered by NetBT name service to a WINS client, even though the servers are running on the same port on the same computer. (Therefore, a machine that runs a WINS server but is not configured as a WINS client won't advertise its own services via WINS!)

20.3.1. Name Resolution Under Windows

For historical reasons, Microsoft provides multiple methods of name resolution. Windows 2000 uses DNS wherever possible, but other icrosoft Windows machines may use numerous methods to translate a name to an address, depending on their configurations. They may try multiple methods to translate the same name, attempting various things until something succeeds.

For this purpose, there are two kinds of names: possible NetBIOS names and non-NetBIOS names. A genuine NetBIOS name is one known to the NetBIOS name service, but a possible NetBIOS name is any name that is up to 15 characters long. (In theory, a name with a period in it is not a possible NetBIOS name, but in fact, it will be treated as one under some versions of Windows NT.) Names that are not possible NetBIOS names (that is, ones that are over 15 characters long) must be resolved via DNS. The methods recent Windows machines can use for possible NetBIOS names are:

- Do a local NetBIOS broadcast and see if a machine by that name answers.
- Check a local file named *lmhosts* or *hosts* for an entry with the relevant name.
- Do a WINS query.
- Do a DNS query.

In addition, they may have the information cached and therefore not need to look it up at all. Windows NT 4 machines that are WINS and DNS clients by default try these options in the following order:

1. Look in the machine's cache of name information; this includes information for machines that have been looked up recently and information from *lmhosts* that has been marked #PRE, and it may include information about hosts on the local network that have booted recently.
2. Do a WINS query.
3. Check a local *lmhosts* file.
4. Do a NetBIOS broadcast.
5. Check a local *hosts* file.
6. Do a DNS query.

Various configuration options will change this order (most notably, if DNS support is not configured, DNS and *hosts* will not be checked). Other operating systems prefer other orders. In particular, older operating systems may try broadcasting before doing a WINS query or may not support any method except broadcasting.

icrosoft has a complicated naming scheme for describing these options; a machine that does only broadcast is a b-node, one that only does WINS queries is a p-node ("p" for "point-to-point"), one that does broadcast and then WINS is an m-node ("m" for "mixed"), and one that does WINS first and then broadcast is an h-node ("h" for "hybrid"). This naming scheme is useful mostly for understanding Microsoft documentation, since knowing what kind of node a machine is does not tell you where in the process it consults DNS or *lmhosts*. However, you may wish to remember that "m-node" and "b-node" both mean a "machine that generates lots of annoying broadcast packets". As you can determine by careful reading of the definitions and the preceding steps, Windows NT machines configured to use WINS are normally h-nodes.

Clients may get DNS data without making DNS queries; WINS servers can act as gateways into DNS, and other machines may cache DNS data and return it in answer to NetBIOS queries. In some configurations, this will result in a puzzling situation where the clients cannot reach hosts with names over 15 characters long. Since these are not valid NetBIOS names, clients must speak DNS directly in order to resolve them.

Just to further confuse matters, Microsoft DNS servers are capable of acting as gateways to WINS servers, making WINS queries to try to resolve DNS queries. This has two advantages; it obviously simplifies DNS administration for sites that are already using WINS, by making it nearly nonexistent. More importantly, WINS supports dynamic host registration. When a host comes up, it tells the WINS server what its name and network address are. By contrast, standard DNS provides no way for a host to give this information to the server; the server must be preconfigured with the name-address mapping. (DNS and dynamic update are discussed earlier.) Forwarding DNS queries to WINS provides dynamic DNS registration without modification of DNS.

The *nbtstat* command will show you information about NetBT names on a machine. Using *nbtstat* to check NetBT naming information and *nslookup* to check DNS naming information will often help you straighten out questions about where information is coming from (and therefore what to fix in order to get things to work correctly).

20.3.2. NetBIOS Names

NetBT name service is used to resolve more than just hostnames. There are two basic types of NetBT names: unique names and group names. A unique name maps to exactly one IP address, while a group name may map to multiple IP addresses. Each of these has multiple subtypes, used for different purposes. For instance, one type of group name is used to indicate a multihomed host; another type is used to indicate an administrative group (all the printers, for instance); and a third type is used for showing domain and workgroup membership. Different types of unique names are used for different services. The sixteenth byte of a NetBIOS name is used to indicate the type, and this byte is typically shown in hexadecimal notation when names are displayed (you may see type 03 displayed as "0x03", "<03>", or "03h").

If you actually look at packet traces, particularly if you are using a packet trace system that is not aware of NetBT, you may see NetBIOS names in mangled form. NetBT is based on DNS packet formats. DNS names are longer than NetBIOS names but can't contain the arbitrary hexadecimal values NetBIOS uses for type indicators. In order to make NetBIOS names fit into DNS name fields, NetBT applies a system that converts each byte of the NetBIOS name into two uppercase ASCII characters, which results in eye-catching, if incomprehensible, name strings. They are made even more noticeable by the fact that the mangling algorithm converts spaces, which are used for padding, into "CA", so that almost every name ends in "CACA". Many English speakers (including presumably the engineer who designed the algorithm) attach no special meaning to this string, but the repeating pattern is still very striking. For details of the name-mangling algorithm, see RFC 1001. If you are on a icrosoft machine, most packet sniffing programs will unmangle the names for you; if you are on a Unix machine, the Samba package contains name mangling and unmangling routines. See Appendix A, "Resources", for information on how to get Samba.

When a computer boots, it registers multiple names of different types. These names are not necessarily based on the hostname of the machine; they are chosen according to the purpose the name will be used for. For instance, one of the NetBT names that machines normally register is a unique name for the Messenger service. This service is used to send messages, and if a user is logged in to the console of the machine, the Messenger service will be registered under the user's name as well as under the computer's name. This is designed to allow you to send messages to people by name (human beings rarely want to talk to computers, after all), but it means that NetBT name registrations will contain not only hostnames, but also usernames, and in many cases will let you figure out whether or not somebody is logged in to a machine.

A machine will also register a group name for the workgroup or domain it is part of. This does not make the machine a valid member of a domain; machines may register as parts of domains they don't belong to, and it will not have any security implications. For clients, there is no difference between the registration for a domain and a workgroup. Domain controllers will make some extra registrations; there is a group name for all of the domain controllers and a unique name for the primary domain controller.

In addition, machines that are running the Windows Browser server will register a number of special names used by this service. These name registrations are sometimes treated specially. See the section on the Windows Browser, later in this chapter, for more information about these names and their registration.

Group names can contain large numbers of hosts, which makes them expensive to maintain. As a result, WINS servers are allowed to skip a lot of the verification they do on hostnames when they're dealing with group names, and invalid group names may remain registered for long periods of time.

20.3.3. NetBT Name Service Operations

DNS is a simple question-and-answer system. NetBT name service, because it's dynamic, is nowhere near as straightforward. Clients and servers interact in numerous ways, in order to register names, refresh and release them, and look up other names. The details of this process are quite intricate, and despite the daunting amount of detail here, a number of special cases and possible interactions have been glossed over.

20.3.3.1. General principles of NetBT operations

There are some generalizations about how NetBT name service works:

- Queries may be broadcast or unicast, but responses are always unicast to the host that made the query.
- WINS servers respond to all requests with positive or negative answers (in fact, if a query takes any significant amount of time to process, they will return an intermediate

answer called a *wait acknowledgment* or WACK). NetBT name servers dealing with broadcast requests, however, answer only if they have something to say.
- Clients will always repeat a query if there is no answer (even if it is a broadcast query to which no answer was expected).

20.3.3.2. Name registration

When a client starts up, it registers the names it wants to answer to, starting with a name registration request. Under NetBT name service, it does this by broadcasting the registration request; with WINS, the request is sent directly to the WINS server. In either case, it may turn out that some other machine already has a name the client wants, which will set off a procedure of name conflict resolution that's discussed later. A WINS server will always answer the request, while NetBT name servers will answer name registration requests only in order to dispute them. If the name is not already taken, the client will send out a name announcement that confirms that it has the name to the same place or places it sent the original request.

For a normally configured Windows NT machine that's trying to register the nonconflicting name "unique" and is configured with a WINS server address, the process looks like this:

1. The machine sends a name registration request for "unique" to the WINS server.
2. The WINS server updates its database and sends a positive response containing a time-to-live (TTL).
3. The machine broadcasts a name registration request for "unique" and gets no response.
4. The machine broadcasts another three name registration requests for "unique", just in case, and still gets no response.
5. The machine sends a name announcement to the broadcast address.
6. All machines that receive the broadcast cache the information temporarily, overwriting any previous entry they had for that name.

20.3.3.3. Name refresh

In order to keep WINS server databases free of old junk data, WINS servers hand out limited-time registrations. Clients need to send a name refresh request before the TTL given to them with the name registration expires. If they don't send a name refresh, the WINS server will eventually remove the name. This process is not particularly fast. Windows NT normally gives out registrations with six-day TTLs, so a client that goes away may still be visible in the database for up to six days if nobody tries to release or register the name.

Clients may also send out name refresh requests for reasons of their own (for instance, a client that changes its configuration to use WINS will send a name refresh request to the WINS server).

20.3.3.4. Name resolution

When a client wants to map a name to an IP address, it sends out a name query request, either by broadcast or to the WINS server. Every machine that has a mapping for the name will reply with that mapping. A WINS server will reply to the request even if it doesn't have the answer. Suppose "unique" wants to talk to "stupid", which is a misconfigured client on the local network that doesn't use WINS and didn't boot recently. The procedure will look like this:

1. "unique" checks its cache but doesn't find "stupid" there (if "stupid" had just booted, it would have worked).
2. "unique" sends a name query request to the WINS server.
3. The WINS server sends a negative response.
4. "unique" broadcasts a name query request.
5. All machines that know where "stupid" is, including "stupid" itself and every machine that's connected to it recently, return responses. "unique" believes the first answer it gets.

20.3.3.5. Name release

When a client shuts down or changes its name, it sends a name release request to the WINS server and/or the broadcast address to give notice that the name is now available. As usual, the WINS server will always respond, while other servers will respond only if they have something to say (in this case, if they reject the name release). A client that is trying to release its own name will consider a name released and continue if it receives any response to the name release request (positive or negative). If it gets no answer, it will try again before continuing.

Name release requests may also be sent by machines other than the client that registered the name. This is supposed to allow a machine to correct invalid data. A machine that receives a response from the server, but gets no response at that address, can send a name release request for the name. The WINS server will then do a name query request and release the name if it fails. Other servers will respond to name release requests only for names that they own, for which they will return negative responses. A client that tries to release somebody else's name pays attention to the contents of the response and does not release a name for which it gets a negative response.

20.3.3.6. Conflict management

What happens if a machine tries to register a name as a unique name, but it's already in use, or tries to refresh a name that some other host has registered? Using broadcasts, when a machine sends out the name registration request, the machine that already has the name registered will respond with a negative response. Machines that have conflicting cached data will not respond.

A WINS server has a somewhat more complicated task. It can't be sure that the existing record it has is for a machine that's still running -- in fact, there's a significant chance that the existing record is actually for the same machine, and it's just been moved from one network to another. Therefore, the WINS server will check to see if the record is correct by sending a name query request for the name to the address it's registered at. If the host is still there, it will answer, and the WINS server will send a negative response to the new registration.

There's an important subtlety to note here. Machines that are not WINS servers normally receive name query requests only via broadcast; unicast name queries are the business of

WINS. But name query requests used for conflict resolution and verification of third-party release requests are unicast, so non-WINS servers must answer them. Machines that don't run name servers cannot defend against conflicting name registrations, even if they otherwise use WINS.

Conflict management is one case where groups have a special exemption. If a host tries to register an existing group name as a unique name, the WINS server is not required to check that the group registration is still valid; it can simply reject the attempt.

20.3.4. WINS Server-Server Communication

As well as picking up information from clients as they boot, WINS servers can also exchange information with other WINS servers. This approach is used to provide redundancy and to allow name information to be propagated across large networks.

Unlike DNS servers, WINS servers have no hierarchical structure and do not normally forward queries from one server to another.[122] WINS servers that talk to each other are trying to cause both servers to have identical databases. The replication protocol uses several tricks to try to distribute only updates rather than entire databases between servers, but there is nothing like a DNS zone that would allow them to subdivide the database.

[122]The protocol allows a WINS server to answer a query by directing the client to query another server, but this facility does not appear to be used in practice; in any case, there is no direct server-server forwarding.

WINS server replication is a complicated topic, involving numerous options (for instance, WINS servers need not replicate symmetrically). The details are beyond the scope of this book but are covered in most references about Windows NT network administration (for instance, in Microsoft's *Windows NT Server Networking Guide*, Microsoft Press, which is part of the Windows NT Server Resource Kit). From a firewall point of view, the interesting points about WINS server replication are:

- WINS servers speak to each other over TCP port 42.
- WINS servers attempt to locate replication partners via multicast and will send out IGMP packets to register multicast addresses for this purpose. IGMP is discussed in Chapter 22, "Administrative Services".
- No matter what kind of WINS replication you establish, a pair of replicating WINS servers will make connections in both directions (you cannot set up replication so that only one of the two servers needs to be able to initiate a connection).

WINS servers may have "push" or "pull" partners, but in fact, data is transferred between them only when the receiving machine requests it. A machine that tries to "push" data will simply inform the other machine that new data is available, and the other machine will then request the data.

By default, WINS servers will send data only to machines that are configured as replication partners. It can be reset via the WINS anager or the registry. It should be left in the default mode to help prevent attackers from pretending to be replication partners and pulling the entire WINS database with all of its information about valid hostnames and usernames.

20.3.5. The WINS Manager

It is possible to control and configure WINS servers on remote machines with the WINS manager, which uses Microsoft RPC. The security implications of icrosoft RPC are discussed in Chapter 14, "Intermediary Protocols".

20.3.6. Security Implications of NetBT Name Service and WINS

NetBT Name Service and WINS are very vulnerable and are much more sensitive than DNS. The information they provide is valuable to attackers; it's not just the hostname data DNS provides, which is already useful, but also information about what usernames are valid and whether anybody is logged in, plus structural information about what machines provide what services. This information gives all sorts of leads to further possible attacks.

In addition, these services modify their databases and take other actions based on information from clients, which vastly magnifies the risks. An attacker who can send packets to a WINS server can cause the WINS server to send packets other places, using it as an amplifier to spread denial of service attacks and as a gateway to get those attacks to networks that may not be directly vulnerable to the attacker. It's also easy for an attacker to contaminate NetBT name service and WINS databases with bad data, which is relatively difficult in DNS.

The protocols are more complex than DNS. The extra protocol layers and headers are all opportunities for implementors to introduce bugs that can become denial of service attacks. Some of these problems have already been found by accident by people implementing these protocols on other platforms and are fixed in recent versions; others are presumably lurking, waiting for malicious people to run out of easier targets.

Finally, DNS servers are relatively rare. Any given site has a handful. NetBT name servers are everywhere. Every machine that uses NetBT name service must also be a NetBT name server. For an attacker who is going after quantity, rather than quality, NetBT name service is a very tempting target; it has information-rich servers everywhere accepting queries and data from arbitrary hosts.

20.3.7. Packet Filtering Characteristics of NetBT Name Service

NetBT name service uses TCP and UDP port 137. Almost all NetBT name service traffic will be UDP; clients normally use TCP only if they issue a query via UDP and get a truncated response because the response is too long to fit into a single UDP packet. However, servers will respond to any query via TCP. Microsoft implementations use port 137 for queries as well as responses. Some older versions will return UDP responses to port 137 regardless of the port the query was made from. Given that most requests are UDP, and both ends are at port 137, it is basically impossible to allow service in a single direction. (Since server-initiated queries are part of conflict resolution, it wouldn't help much anyway.)

Clients that do not use WINS will send queries to the broadcast address; responses always are unicast, as are WINS queries. WINS servers will try to use multicast to contact replication partners and will therefore generate IGMP packets (see Chapter 22, "Administrative Services", for packet filtering details of IGMP). WINS server replication uses TCP port 42.

Direction	SourceAddr.	Dest.Addr.	Protocol	SourcePort	Dest.Port	ACKSet	Notes
In	Ext	Broadcast	UDP	137, >1023	137	[123]	Incoming NetBT name service query via UDP, client to server
In	Ext	Int	UDP	137, >1023	137	[123]	Incoming WINS query via UDP, client to server
Out	Int	Ext	UDP	137	137, >1023		Answer to incoming UDP query, server to client
In	Ext	Int	TCP	137, >1023	137	[124]	Incoming query via TCP, client to server
Out	Int	Ext	TCP	137	137, >1023	Yes	Answer to incoming TCP query, server to client
Out	Int	Broadcast	UDP	137, >1023	137	[123]	Outgoing NetBT name service query via UDP
Out	Int	Ext	UDP	137, >1023	137	[123]	Outgoing WINS query via UDP
In	Ext	Int	UDP	137	137, >1023	[123]	Answer to outgoing UDP query
Out	Int	Ext	TCP	137, >1023	137	[124]	Outgoing query via TCP, client to server
In	Ext	Int	TCP	137	137,	Yes	Answer to

					>1023		outgoing TCP query, server to client
Out	Int	Ext	TCP	>1023	42	[124]	WINS server replication request from internal server to external server
In	Ext	Int	TCP	42	>1023	Yes	WINS server replication reply
In	Ext	Int	TCP	>1023	42	[124]	WINS server replication request from external server to internal server
Out	Int	Ext	TCP	42	>1023	Yes	WINS server replication reply

[123]UDP has no ACK equivalent.

[124]ACK is not set on the first packet of this type (establishing connection) but will be set on the rest.

20.3.8. Proxying Characteristics of NetBT Name Service and WINS

icrosoft provides something called a *WINS proxy service*, which turns broadcast NetBIOS name resolution requests into WINS requests, in order to allow non-WINS clients to use a WINS server. This is a form of proxying, although it also is changing the protocol in use.

It would be perfectly possible to do more traditional firewall proxying of WINS, but there do not appear to be any implementations available.

20.3.9. Network Address Translation Characteristics of NetBT Name Service and WINS

Because these protocols are attempting to maintain mappings between names and IP addresses, they frequently contain embedded IP addresses. Furthermore, these addresses are often deep in the content of the packet (not simply in the NetBT destination headers used by many other NetBT-based protocols). A network address translator would have to be aware of the details of the protocol in order to successfully translate all of the embedded addresses. In any case, you cannot save address space by using a network address translator if you run the name service through it, since all hosts will attempt to register their names and addresses.

20.3.10. Summary of Recommendations for NetBT Name Service and WINS

- Do not allow WINS queries or server replication across your firewall.
- Bastion hosts that are configured not to respond to name requests will be unable to defend themselves from other hosts that try to take over their names via NetBT. These machines should be statically configured into WINS servers or accessed via 16-character names that cannot be resolved with NetBT. See Chapter 10, "Bastion Hosts", for more information.

How to Configure Network Load Balancing for Configuration Manager Site Systems

Updated: July 1, 2010

Applies To: System Center Configuration Manager 2007, System Center Configuration Manager 2007 SP1, System Center Configuration Manager 2007 SP2

Network Load Balancing (NLB) clusters provide scalability in Configuration Manager 2007 so that you can support more than 25,000 clients at one site.

Windows Server Network Load Balancing distributes client requests across a set of servers and supports up to 32 computers running Windows Server in a single cluster. When Network Load Balancing is installed as a network driver on each of the member servers (hosts) in a cluster, the cluster presents a virtual IP address or fully qualified domain name (FQDN) to client requests. The initial client requests go to all the hosts in the cluster, but only one host accepts and handles the request.

All service pack levels of Configuration Manager 2007 support using NLB clusters for the following site system roles:

- Management point
- Software update point
- Server locator point

Each host in a NLB cluster must meet the supported configuration requirements for site systems of the Configuration Manager service pack level in use. Configuration Manager 2007 supports use of NLB clusters on the following operating systems:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2003

- Windows Server 2003 R2

Use the following information on this page to configure NLB clusters for Configuration Manager:

- Planning for NLB Clusters in Configuration Manager
- Configure Windows Server Computers as Members of a NLB Cluster
- Install the Site System Role on NLB Cluster Members
- Configure NLB Management Points in Mixed Mode Sites for Client Approval
- Configure NLB Management Points and NLB Software Update Points in Native Mode Sites
- Designate the NLB Cluster for Configure Configuration Manager

Planning for NLB Clusters in Configuration Manager

All computers that will be part of a NLB cluster for Configuration Manager have the following requirements:

- All computers in the NLB cluster must be in the same domain.
- Each computer in the NLB cluster must use a static IP address.
- Each computer in the NLB cluster must have **Network Load Balancing** enabled.
- In a native mode site, the NLB cluster must be configured for a FQDN.
- The NLB cluster requires a static IP address.

Configure Windows Server Computers as Members of a NLB Cluster

Although the exact steps to configure a Windows Server as part of a NLB cluster depend upon the Windows Server version in use, all versions have the following configuration requirements:

- The **Cluster operation** mode must be set to **Unicast**.
- In a Configuration Manager 2007 native mode site, the NLB cluster requires a FQDN.
- You must manually register the NLB cluster name in DNS by using a host (A) or (AAAA) record because DNS does not automatically register static IP addresses.

To configure the NLB cluster for Configuration Manager, see the following guidance for the Windows Server operating system in use.

- To configure NLB clusters on Windows Server 2008, see Creating Network Load Balancing Clusters (http://go.microsoft.com/fwlink/?LinkId=197176) in the Windows Server 2008 TechNet library.
- To configure NLB clusters on Windows Server 2008 R2, see Creating Network Load Balancing Clusters (http://go.microsoft.com/fwlink/?LinkId=197177) in the Windows Server 2008 R2 TechNet library.
- To configure NLB clusters on Windows Server 2003 and Windows Server 2003 R2, use the following procedure to implement Network Load Balancing for Windows Server 2003 and Windows Server 2003 R2 Configuration Manager 2007 site systems.

To configure Network Load Balancing for Configuration Manager site system computers using NLB.exe on Windows Server 2003 and Windows Server 2003 R2

1. Click **Start**, click **All Programs**, click **Administrative Tools**, and then click **Network Load Balancing Manager**.

2. On the menu bar, click **Cluster**, and then click **New** to open the **Cluster Parameters** dialog box.

3. In the **Cluster Parameters** dialog box, enter the information for the Network Load Balancing cluster IP configuration:

   - **IP address:** NLB cluster IP address registered in DNS

   - **Subnet mask:** NLB cluster IP address subnet mask registered in DNS

   - **Full Internet name:** FQDN of NLB cluster name registered in DNS

4. Ensure **Unicast** is selected in **Cluster operation mode** and then click **Next**.

5. On the **Cluster IP Addresses** page, click **Next**.

6. On the **Port Rules** page, click **Edit** to define the ports that the NLB cluster will respond to, and configure the ports used for client to site system communication defined for the site, or click **Next** to enable the NLB cluster IP address to respond to all TCP/IP ports.

   > **Note**
   >
   > Ensure that **Affinity** is set to **Single**.

7. On the **Connect** page, enter a site system host name that will be part of the NLB cluster in **Host**, and then click **Connect**.

8. In **Interfaces available for configuring a new cluster**, select the networking interface that will be configured to respond to NLB cluster communication, and then click **Next**.

9. On the **Host Parameters** page, review the information displayed to ensure that the **Dedicated IP configuration** settings display the dedicated host IP configuration for the correct NLB cluster host, the Initial host state **Default state:** is **Started**, and then click **Finish**.

   > **Note**
   >
   > The **Host Parameters** page also displays the NLB cluster host priority (1-32). As new hosts are added to the NLB cluster, the host priority must differ from the previously added hosts. The priority is automatically incremented when using the Network Load Balancing Manager.

10. Click **<NLB cluster name>** and ensure that the NLB host interface **Status** displays **Converged** before continuing. This step might require refreshing the NLB cluster display as the host TCP/IP configuration is being modified by Network Load Balancing Manager.

11. To add additional hosts to the NLB cluster, right-click **<NLB cluster name>**, click **Add Host to Cluster,** and then repeat steps 7 through 10 for each site system that will be part of the NLB cluster.

Install the Site System Role on NLB Cluster Members

After you have configured the NLB cluster, you must install the Configuration Manager site system role on each computer (host) in the NLB cluster. Ensure that the site system roles are installed and functioning correctly before you continue.

For more information about adding site system roles to Configuration Manager 2007 site systems, see How to Add New Site System Roles.

Configure NLB Management Points in Mixed Mode Sites for Client Approval

When your site is operating in mixed mode, and you are using automatic client approval, the following additional configuration is required:

- Create an Internet Information Services (IIS) application pool service account in Active Directory Domain Services and register a Service Principal Name (SPN) for the account.
- Configure this account to run the **CCM Windows Auth Server Framework Pool** in Internet Information Services (IIS) Manager on each management point that is configured as part of the Network Load Balancing cluster.

> **Important**
>
> This additional configuration is unnecessary for sites configured to operate in native mode. In native mode, use SSL certificates for client authentication. Configuring an application pool service account SPN in Active Directory Domain Services is not required.

For more information about configuring and registering an SPN for the IIS application pool for Configuration Manager management point site systems configured in an NLB cluster, see How to Configure an SPN for NLB Management Point Site Systems.

After the SPN for the IIS application pool is registered in Active Directory Domain Services, you must configure each computer in the NLB cluster to use this account.

To configure the IIS application pool service account for management point site systems configured in Network Load Balancing clusters

1. Click **Start**, click **All Programs**, click **Administrative Tools**, and then click **Internet Information Services (IIS) Manager** to open the Internet Information Services (IIS) Manager console.
2. Expand <*computer name*>, expand **Application Pools**, right-click **CCM Windows Auth Server Framework Pool**, and then click **Properties**.
3. In the **CCM Windows Auth Server Framework Pool Properties**, click the **Identity** tab.
4. Select **Configurable** and enter the application pool service account **User name** and **Password** information, and then click **OK**. In the **Confirm Password** dialog box, confirm the application pool service account password.

Configure NLB Management Points and NLB Software Update Points in Native Mode Sites

When you configure Network Load Balancing for Configuration Manager management points or software update points for sites that are in native mode, verify that the certificates are configured correctly. Each member server in the NLB cluster must have a public key infrastructure (PKI) certificate that contains both the FQDN of the NLB site system and the site system server name. For more information, refer to the certificate requirements listed in the section "Network Load Balancing Management Points or Network Load Balancing Software Update Points" in Certificate Requirements for Native Mode.

Designate the NLB Cluster for Configuration Manager

You must specify the Network Load Balanced management points and software update points in the Configuration Manager console.

> **Important**
>
> For management points on NLB clusters in mixed mode sites, you must complete the configuration of the IIS application pool on each NLB cluster member before you proceed with the following configuration.

To specify the NLB management points and software update points in Configuration Manager

1. In the Configuration Manager console, navigate to **System Center Configuration Manager** / **Site Database** / **Site Management** / **<site code> - <site name>** / **Site Settings** / **Component Configuration**.

2. Right-click either **Management Point Component** or **Software Update Point Component**, and then click **Properties**.

3. On the **General** tab, select the option to use Network Load Balancing clusters and then configure the IP address or FQDN of the virtual server hosting the Network Load Balancing cluster.

Server locator point site systems configured as Network Load Balancing clusters are not defined in the Configuration Manager console. Instead, specify the server locator point NLB during client installation by using the **/SMSSLP=<server locator point NLB cluster name>** command-line option.

If you use WINS, you must manually add the site system information to WINS. For information about manually adding Configuration Manager site information to WINS, see How to Manually Add Configuration Manager Site Information to WINS.

See Also

Tasks

How to Configure an SPN for NLB Management Point Site Systems
How to Configure the Default Management Point for a Site
How to Configure the Intranet FQDN of an NLB Management Point
How to Configure HTTP Communication for Roaming and Site Assignment

Concepts

Configuration Manager and Network Load Balancing Clusters
Configuration Manager and Service Location (Site Information and Management Points)

Configuring Network Interfaces and the Boot Device

This chapter provides information and instructions that are required to plan and to configure the supported network interfaces.

Tasks covered in this chapter include:

- How to Configure the Primary Network Interface
- How to Configure Additional Network Interfaces
- How to Select the Boot Device

Note - Many of the procedures in this chapter assume that you are familiar with the OpenBoot firmware and that you know how to enter the OpenBoot environment. For background information, refer to About the ok Prompt. For instructions, refer to How to Get to the ok Prompt.

⚠ **Caution -** Do not attempt to access any internal components unless you are a qualified service technician. Detailed service instructions can be found in the *Sun Fire V490 Server Parts Installation and Removal Guide*, which is included on the Sun Fire V490 Documentation CD.

How to Configure the Primary Network Interface

Before You Begin

You must perform this task:

- Complete the installation steps in Chapter 1.

For background information, refer to:

- About the Network Interfaces

If you are using a PCI network interface card, refer to the documentation supplied with the card.

What to Do

1. **Choose a network port, using the following table as a guide.**

Ethernet Port	PCI Bus/Clock Rate	OpenBoot devalias	Device Path
1	PCI C/66 MHz	net1	/pci@9,600000/network@1
0	PCI D/33 MHz	net0	/pci@9,700000/network@2

2. Attach an Ethernet cable to the port you selected.

Refer to How to Attach a Twisted-Pair Ethernet Cable.

3. Choose a host name for the system and make a note of it.

You need to furnish the name in a later step.

The host name must be unique within the network. It can consist only of alphanumeric characters and the dash (-). Do not use a dot in the host name. Do not begin the name with a number or a special character. The name must not be longer than 30 characters.

4. Determine the unique Internet Protocol (IP) address of the network interface and make a note of it.

You need to furnish the address in a later step.

An IP address must be assigned by the network administrator. Each network device or interface must have a unique IP address.

5. Resume the installation of the system.

Return to Chapter 1.

Note - During installation of the Solaris OS, the software automatically detects the system's on-board network interfaces and any installed PCI network interface cards for which native Solaris device drivers exist. The operating system then asks you to select one of the interfaces as the primary network interface and prompts you for its host name and IP address. You can configure only one network interface during installation of the operating system. You must configure any additional interfaces separately, after the operating system is installed. For more information, refer to How to Configure Additional Network Interfaces.

What Next

After completing this procedure, the primary network interface is ready for operation. However, in order for other network devices to communicate with the system, you must enter the system's IP address and host name into the namespace on the network name server. For information about setting up a network name service, consult:

- *Solaris Naming Configuration Guide* for your specific Solaris release

The device driver for the system's on-board Sun GigaSwift Ethernet interfaces is automatically installed with the Solaris release. For information about operating characteristics and configuration parameters for this driver, refer to the following document:

- *Platform Notes: The Sun GigaSwift Ethernet Device Driver*

This document is available on the Solaris Software Supplement CD for your specific Solaris release.

If you want to set up an additional network interface, you must configure it separately, after installing the operating system. Refer to:

- How to Configure Additional Network Interfaces

Note - The Sun Fire V490 system conforms to the Ethernet 10/100BASE-T standard, which states that the Ethernet 10BASE-T link integrity test function should always be enabled on both the host system and the Ethernet hub. If you have problems establishing a connection between this system and your hub, verify that the Ethernet hub also has the link test function enabled. Consult the manual provided with your hub for more information about the link integrity test function.

How to Configure Additional Network Interfaces

Before You Begin

Perform the following tasks to prepare an additional network interface:

- Install the Sun Fire V490 server as described in Chapter 1.
- If you are setting up a redundant network interface, refer to About Redundant Network Interfaces.
- If you need to install a PCI network interface card, follow the installation instructions in the *Sun Fire V490 Server Parts Installation and Removal Guide*.
- Attach an Ethernet cable to the appropriate port on the system back panel. Refer to How to Attach a Twisted-Pair Ethernet Cable. If you are using a PCI network interface card, refer to the documentation supplied with the card.

Note - All internal options (except disk drives and power supplies) must be installed by qualified service personnel. Installation procedures for these components are covered in the *Sun Fire V490 Server Parts Installation and Removal Guide*, which is included on the Sun Fire V490 Documentation CD.

What to Do

1. Choose a network host name for each new interface.

The host name must be unique within the network. It can consist only of alphanumeric characters and the dash (-). Do not use a dot in the host name. Do not begin the name with a number or a special character. The name must not be longer than 30 characters.

Usually an interface host name is based on the machine host name. For example, if the machine is assigned the host name `sunrise`, the added network interface could be named `sunrise-1`. The machine's host name is assigned when Solaris software is installed. For more information, refer to the installation instructions accompanying the Solaris software.

2. Determine the Internet Protocol (IP) address for each new interface.

An IP address must be assigned by your network administrator. Each interface on a network must have a unique IP address.

3. Boot the operating system (if it is not already running) and log on to the system as superuser.

Be sure to perform a reconfiguration boot if you just added a new PCI network interface card. Refer to How to Initiate a Reconfiguration Boot.

Type the `su` command at the system prompt, followed by the superuser password:

```
% su
Password:
```

4. Create an appropriate `/etc/hostname` file for each new network interface.

The name of the file you create should be of the form `/etc/hostname.ce`*num*, where `ce` is the network interface type identifier and *num* is the device instance number of the interface according to the order in which it was installed in the system.

For example, the file names for the system's on-board Sun GigaSwift Ethernet interfaces are `/etc/hostname.ce0` and `/etc/hostname.ce1`, respectively. If you add a PCI Ethernet adapter card as a third `ce` interface, its file name should be `/etc/hostname.ce2`. At least one of these files--the primary network interface--should exist already, having been created automatically during the Solaris installation process.

Note - The documentation accompanying the network interface card should identify its type. Alternatively, you can enter the `show-devs` command from the `ok` prompt to obtain a list of all installed devices.

5. Edit the `/etc/hostname` file(s) created in Step 4 to add the host name(s) determined in Step 1.

Following is an example of the `/etc/hostname` files required for a system called `sunrise`, which has two on-board Sun GigaSwift Ethernet interfaces (`ce0` and `ce1`) and a PCI Ethernet adapter card (`ce2`). A network connected to the on-board `ce0` and `ce1` interfaces will know the system as `sunrise` and `sunrise-1`, while networks connected to the PCI-based `ce2` interface will know the system as `sunrise-2`.

```
sunrise #  cat /etc/hostname.ce0
sunrise
sunrise #  cat /etc/hostname.ce1
sunrise-1
sunrise #  cat /etc/hostname.ce2
sunrise-2
```

6. Create an entry in the `/etc/hosts` file for each active network interface.

An entry consists of the IP address and the host name for each interface.

The following example shows an `/etc/hosts` file with entries for the three network interfaces used as examples in this procedure.

```
sunrise #  cat /etc/hosts
#
# Internet host table
#
127.0.0.1     localhost
129.144.10.57 sunrise loghost
129.144.14.26 sunrise-1
129.144.11.83 sunrise-2
```

7. Manually plumb and enable each new interface using the `ifconfig` command.

For example, for the interface `ce2`, type:

```
sunrise # ifconfig ce2 plumb up
```

For more information, refer to the `ifconfig(1M)` man page.

What Next

After completing this procedure, any new network interfaces are ready for operation. However, in order for other network devices to communicate with the system through the new interface, the IP address and host name for each new interface must be entered into the namespace on the network name server. For information about setting up a network name service, consult:

- *Solaris Naming Configuration Guide* for your specific Solaris release

The `ce` device driver for the system's on-board Sun GigaSwift Ethernet interfaces is automatically configured during Solaris installation. For information about operating characteristics and configuration parameters for these drivers, refer to

- *Platform Notes: The Sun GigaSwift Ethernet Device Driver*

This document is available on the Solaris Software Supplement CD for your specific Solaris release.

Note - The Sun Fire V490 system conforms to the Ethernet 10/100BASE-T standard, which states that the Ethernet 10BASE-T link integrity test function should always be enabled on both the host system and the Ethernet hub. If you have problems establishing a connection between this system and your Ethernet hub, verify that the hub also has the link test function enabled. Consult the manual provided with your hub for more information about the link integrity test function.

How to Select the Boot Device

The boot device is specified by the setting of an OpenBoot firmware configuration parameter called `boot-device`. The default setting of this parameter is `disk`

net. Because of this setting, the firmware first attempts to boot from the system hard drive, and if that fails, from the on-board Sun GigaSwift Ethernet interface.Before You Begin

Before you can select a boot device, you must complete system installation according to the instructions in Chapter 1.

Specifically, you must set up a system console and power on the system. Refer to:

- How to Set Up an Alphanumeric Terminal as the System Console
- How to Configure a Local Graphics Terminal as the System Console
- How to Power On the System

If you want to boot from a network, you must also connect the network interface to the network and configure the network interfaces. Refer to:

- How to Attach a Twisted-Pair Ethernet Cable
- How to Configure the Primary Network Interface
- How to Configure Additional Network Interfaces

What to Do

This procedure assumes that you are familiar with the OpenBoot firmware and that you know how to enter the OpenBoot environment. For more information, refer to About the ok Prompt.

- **At the ok prompt, type:**

```
ok setenv boot-device device-specifier
```

where the *device-specifier* is one of the following:

- cdrom - Specifies the DVD-ROM drive
- disk - Specifies the system boot disk
- disk0 - Specifies internal disk 0
- disk1 - Specifies internal disk 1
- net, net0, net1- Specifies the network interfaces
- *full path name* - Specifies the device or network interface by its full path name

Note - You can also specify the name of the program to be booted as well as the way the boot program operates. For more information, refer to the *OpenBoot 4.x Command Reference Manual*, included with the Solaris Software Supplement CD that ships with Solaris software.

If you want to specify a network interface other than an on-board Ethernet interface as the default boot device, you can determine the full path name of each interface by typing:

8.10. Configuring and using a tuntap network interface

If you use linux (optionally FreeBSD and Solaris, not tested), you may want to access the network through a tuntap interface. The main advantage of this interface, is that the guest has access to the host. The guest can even have access to the whole network if the host routes or masquerades the guest requests. No extra IP address is needed, all can be done using private IP addresses.

You'll find here instructions to set up Linux/Bochs to provide network access to the guest OS through a tuntap interface and private IP network. We're going to see howto :

- enable the tuntap interface in the Linux Kernel
- configure Bochs to use the tuntap interface
- set up the private network between the host and the guest
- set up the host to masquerade the guest network accesses

8.10.1. Tuntap description

From the tuntap.txt file in the Linux kernel tree :

```
  TUN/TAP provides packet reception and transmission for user space
programs.
  It can be viewed as a simple Point-to-Point or Ethernet device, which
  instead of receiving packets from a physical media, receives them from
  user space program and instead of sending packets via physical media
  writes them to the user space program.

  When a program opens /dev/net/tun, driver creates and registers
corresponding
  net device tunX or tapX. After a program closed above devices, driver
will
  automatically delete tunXX or tapXX device and all routes corresponding
to it.
```

8.10.2. Set up the linux Kernel [1]

First make sure the tuntap module is included in the kernel :

- if you use a recent distribution, chances are that the needed modules are already build

  Make sure that "Kernel module loader" - module auto-loading support is enabled in your kernel.

  Add following line to the /etc/modules.conf:

  ```
   alias char-major-10-200 tun
  ```

  Run:

  ```
   depmod -a
  ```

  The driver will be automatically loaded when application access /dev/net/tun.

- Otherwise, recompile the kernel, including the configuration option

  ```
   CONFIG_TUN (Network device support -> Universal TUN/TAP device
  driver support)
  ```

Note: Make sure there is a /dev/net/tun device. (Can be created with '**mkdir /dev/net ; mknod /dev/net/tun c 10 200**').

In the same way, to use masquerading, you need a kernel with the following options :

```
 CONFIG_IP_NF_CONNTRACK (Connection tracking)
 CONFIG_IP_NF_IPTABLES (IP tables support)
 CONFIG_IP_NF_NAT (Full NAT)
```

Note: Some of the other options in this group is probably also needed, (but the default setting should be OK).

8.10.3. Configure Bochs to use the tuntap interface

Make sure Bochs has ne2000 support. If you have to recompile Bochs, **--enable-ne2000** when running **./configure** (see Section 3.4)

edit your *.bochsrc* configuration file and add something like :

```
  ne2k: ioaddr=0x300, irq=9, mac=fe:fd:00:00:00:01,
                          ethmod=tuntap, ethdev=/dev/net/tun0,
script=/path/to/tunconfig
```

Since the tuntap interface cannot be configured until a process opens it, Bochs may run a script file for you. In this case */path/to/tunconfig* should be changed to match the actual place where you'll create this script.

8.10.4. Set up the private network between the host and the guest

We'll set up a private network between the host and the guest with the following parameters:

```
                Host IP : 192.168.1.1
                Guest IP : 192.168.1.2
```

If your parameters are different, adapt the rest of the section to suit your needs.

Create the *`/path/to/tunconfig`* script :

```
        #!/bin/bash
        /sbin/ifconfig ${1##/*/} 192.168.1.1
```

The script get the interface name as the first parameter. Linux will forward incoming packets between interfaces.

Make it executable :

```
chmod 755 /path/to/tunconfig
```

Run Bochs, install the guest OS, and set the following network parameters in the guest OS:

```
        IP: 192.168.1.2
        netmask: 255.255.255.0
        gateway: 192.168.1.1
        nameserver: whatever is used in linux
```

Note: Bochs must be started by root (at least for now - the script won't have root privileges otherwise).

You may also have to edit /etc/hosts.allow in the host OS and add :

```
        ALL: 192.168.1.2
```

Don't forget to set up the route on the guest.

At this point, you should be able to ping/telnet/ftp/ssh the guest from the host and vice-versa.

8.10.5. Set up the host to masquerade the guest network accesses

We are going to set up standard masquerading configuration. Edit the *`/path/to/tunconfig`* script ans add :

```
        /sbin/iptables -D POSTROUTING -t nat -s 192.168.1.0/24 -d !
192.168.1.0/24 -j MASQUERADE >& /dev/null
```

```
        /sbin/iptables -t nat -s 192.168.1.0/24 -d ! 192.168.1.0/24 -A
POSTROUTING -j MASQUERADE
        echo 1 > /proc/sys/net/ipv4/ip_forward
```

Note: The configuration assumes the default policy is ACCEPT (can be examined by doing '**/sbin/iptables -L**')

Note: The iptables package must be installed.

Managing DHCP, Windows Internet Name Service, and Internet Authentication Service

A network administrator might use SNMP to assist in the following duties:

- Viewing and changing parameters in the LAN Manager and MIB-II MIBs.
- Monitoring and configuring parameters for any WINS servers on the network.
- Monitoring DHCP servers.
- Using System Monitor to monitor TCP/IP- related performance counters (Internet Control Message Protocol (ICMP), IP, Network Interface, TCP, UDP, DHCP, FTP, WINS, and IIS performance counters).

For more information about System Monitor, see the Microsoft® Windows® 2000 Professional Resource Kit.

Use the tools on the *Windows 2000 Resource Kit* companion CD to perform simple SNMP management functions.

Using System Monitor

All System Monitor counters installed on a computer can be viewed with SNMP. To view System Monitor counters with SNMP, use the Perf2MIB tool provided on the *Windows 2000 Resource Kit* companion CD. For additional information about how to use the Perf2mib.exe tool, see Tools Help on the companion CD.

Top Of Page

Managing DHCP

The Windows 2000 – based DHCP server objects and IIS objects can be monitored but not configured by using SNMP .

Top Of Page

Managing WINS

All but a few of the WINS server objects can be monitored and configured by using SNMP. For information about what WINS parameters can be configured using SNMP, see "MIB Object Types" in this book. Any WINS objects defined with read/write permissions can be configured.

Top Of Page

Managing IAS

Internet Authentication Server (IAS) implements the RADIUS authentication and accounting MIBs, which permit IAS objects to be monitored and configured using SNMP. Any IAS objects defined with read/write permissions can be configured.

TRIP Tutorial

Introduction

TRIP (Telephony Routing over IP, RFC 3219) is a policy driven dynamic routing protocol for advertising the reachability of telephony destinations, and for advertising attributes of the routes to those destinations. TRIPs operation is independent of any signaling protocol, hence TRIP can serve as the telephony routing protocol for any signaling protocol.

TRIP is a routing protocol based on a well-known Internet routing protocol called BGP-4 that forms much of backbone of the Internet today. Like BGP-4, TRIP is a protocol that helps in exchange of routing information among various "TRIP Speakers". TRIP Speakers, also called Location Servers, exchange and store routing information for reachibility of telephony prefixes. Session Protocols like H.323 can then query the Location Server for routes to reach a particular telephony prefix. It is possible that the session protocol might be located on an entity, typically a gateway or gatekeeper in case of H.323, while the Location Server may be running on a separate entity. In this case a RAS like protocol can query the Location Server (analogous to an LRQ-LCF exchange).

TRIP introduces a concept of IP Telephony Administrative Domains (ITADs). An ITAD is a set of resources consisting of gateways and at least one TRIP Location Server (among other VoIP elements) that are controlled by a single authority. In an H.323 Network, an ITAD could consists of a set of H.323 gateways interested in advertising prefixes via the TRIP Speaker. Gateways interested in advertising the prefixes they terminate can "register" with the TRIP Speaker. Protocol for such registration is not specified. A RAS like protocol can be used for such purposes.

An ITAD can span multiple H.323 zones. On the other hand, an ITAD may or may not have the same boundaries as that of an Annex-G domain. A TRIP speaker communicates with other TRIP speakers, both within the same ITAD (intra-domain) and to speakers outside the ITAD (inter-domain).

TRIP Peers

A TRIP Speaker establishes intra-domain and inter-domain "peering sessions" with other TRIP Speakers to exchange routing information. The peering sessions are established to exchange routes to telephony destinations. Two peers update each other of new routes they learn. Each peer may in-turn learn about new routes from other peers or through gateways registering telephony prefixes to them or through a static configuration on the Location Servers. The peers also "withdraw" the routes they advertised to the other peer on learning about the unavailability of the routes. TRIP does not require periodic refresh of the routes. The speakers periodically exchange keep-alive messages to ensure that the peers are operational.

TRIP peering sessions use TCP for transport. Thus the reliability of information is ensured.

Apart from conveying the telephony destinations (prefixes) that a Location Server can reach, a routing update also carries some more information about that route, called the "attributes" associated with the route. As will be discussed later, these attributes are helpful in describing characteristics of the route as well as in correct operation of the protocol. They also help in enforcing policies and network design.

TRIP qualifies inter-domain sessions as running E-TRIP sessions (External TRIP) and intra-domain sessions as I-TRIP (internal TRIP).Figure 1 shows two ITADs. ITAD 1 has two Location Servers. Gateways G1 and G2 register with LS2 and Gateways G3 and G4 register with LS1. LS1 and LS2 have I-TRIP peering. LS1 peers with LS3 in ITAD2 (E-TRIP peering).

Internal TRIP uses a link state mechanism to flood database updates over an arbitrary topology. An attempt is made to synchronize routing information among TRIP LSs within an ITAD to maintain a single unified view. To achieve internal synchronization, internal peer connections are configured between LSs of the same ITAD such that the resulting intra-domain Location Server topology is connected and sufficiently redundant. When an update is received from an internal peer, the routes in the update are checked to determine if they are newer than the version already in the database. Newer routes are then flooded to all other peers in the same ITAD.

While updates within an ITAD are flooded onto internal peers, External TRIP updates are point-to-point.

TRIP updates received by an ITAD X from ITAD Y can be passed on to ITAD Z with or without any modifications (with X and Z not sharing any peering relation). Thus a route "advertisement" might reach a peer after hopping through various TRIP Speakers in different ITADs. The list of ITADs through which an update traverses for a given route is recorded in an attribute associated with that route, called the AdvertismentPath. Each route also has a "NextHopServer" attribute associated with it. It signifies the address of the next signaling entity that needs to be contacted in order to setup a call. It is upto the LocationServer at the boundary of the ITAD if it wants to add a new signalling element in the signaling path for the route in question. For example, an LS could change the NextHopServer for all the routes that it is disseminating to point to a gatekeeper so that all the calls for destination prefixes in that ITAD land on the gatekeeper for billing purposes. In order to do this, the LS would overwrite the NextHopServer of the route (with the address of this new signaling element) before sending out the update to its peers. Thus, a signaling message (like H.225 Setup) for a given telephony destination shall only go through ITADs which have changed the NextHopServer attribute for the route in question. This list of ITADs is recorded in another attribute called RoutedPath.

A TRIP Speaker on the boundary of the ITAD may optionally "aggregate" some routes before sending out a routing update to an external speaker. In Figure 1 we see LS1 aggregating routes from G1 (8581* through 8587*) and G2 (8588* through 8580*) before sending the ETRIP update to LS3 as a 858* route.

Thus TRIP can be used for inter-domain as well as intra-domain routing. It is also possible to use TRIP on a gateway as a registration protocol. When used in this way, the TRIP Protocol shall run on the gateway in a "send-only" mode, only sending routing information (prefixes supported by the gateway) to it's peer (a Location Server).

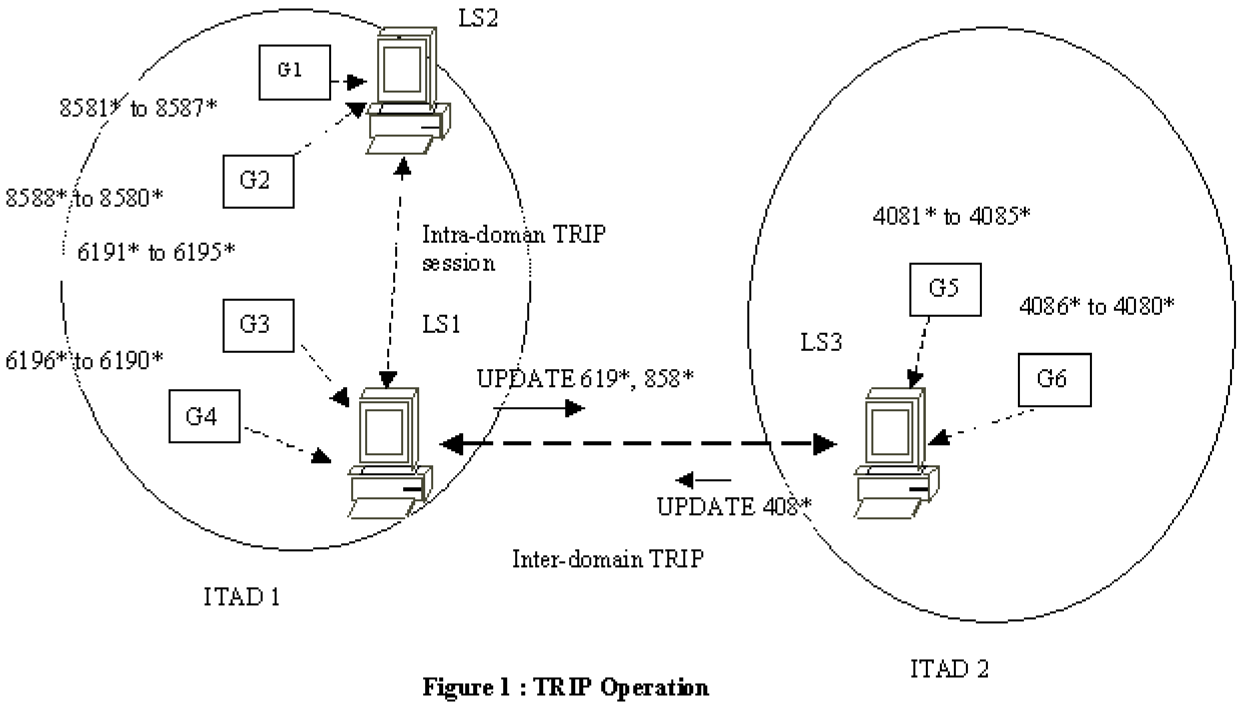

Figure 1 : TRIP Operation

Routes and Route Selection

A route in TRIP is defined as the combination of (a) a telephony destination addresses, given by an address family indicator and an address (telephony) prefix and (b) an application protocol (e.g H.323, SIP etc.). The transport address associated with this telephony destination is included in the "NextHopServer" attribute of the route. TRIP supports two address families: POTS numbers and Routing numbers. POTS numbers address family is a super set of all E.164 numbers, national numbers, local numbers and private numbers. The type of POTS number (private, local, national or international) can be deduced from the first few digits of the POTS prefix. Routing Numbers family is used to represent routing numbers used in conjunction with Number Portability. Routing Numbers are used to identify switches and line cards in the switches. It is possible to extend TRIP to support new address families if needed.

As mentioned earlier, TRIP is a general routing protocol that can be used for disseminating reachibility information of telephony destinations, irrespective of the application (signaling) protocol in use. The TRIP Location Server can be queried to fetch a route for a particular telephony prefix and application protocol combination.TRIP supports four application protocols: SIP, H.323-Q.931, H.323-RAS, and H.323-Annex-G. Thus TRIP can help in propagation of routes even when RAS or H.323-AnnexG is used (For illustrative examples, please refer to the section "TRIP in an H.323 Network")

A Route Selection algorithm runs on the Location Server. The purpose of this algorithm is to generate the best route for a given prefix (and application protocol) based on the information stored in the database. LS returns this as the route for a prefix when queried by an application protocol. The route thus obtained is also the one to be advertised by the LS to its peers, subject to policy configuration. The algorithm comes into play whenever some routing information changes because of new routing updates, introducing or withdrawing routes to certain prefixes. It bases it's decision on certain attributes associated with the routes that define the characteristics of the path associated with the route.

The notion of attributes in TRIP plays an important role in correct and efficient functioning of the protocol. The RoutedPath attribute is used to specify the intermediate ITADs to be taken by the signaling protocol to reach the destination prefix. RoutedPath may be used in selection of a route when multiple routes are present: An LS may prefer route with a lower number of ITADs in the RoutedPath attribute so that a signaling call has to go through minimal hops. An Advertisement Path traces the ITADs that a route advertisement (update) has traveled so far. AdvertisementPath is useful in detecting loops in the routes. Local Preference attribute helps in determining a particular LS's preference for a route. This can help divert intra-domain signaling traffic to a particular ETRIP peer when more than one route for a telephony destination is available. Diversion of traffic might be helpful if signaling elements in a certain neighboring ITAD have higher capacity to handle signaling traffic. TRIP also supports load balancing of traffic across two or more links between two ITADs. This is achieved by the use of an attribute called MultiExitDisc. Other attributes like Atomic Aggregate and Communities are also helpful in facilitating efficient routing. Thus attribute information can be very helpful in traffic engineering of networks.

The TRIP specification lends itself to extensibility by defining new attributes. For example attributes for defining cost and capacity could be added in the future. Codes for the new attributes can be obtained from IANA.

As mentioned earlier TRIP, like BGP, is a policy driven protocol. For a peering session, it is possible to define filters for incoming and outgoing routing information. For example, a Location Server may want to accept a routing update only if the attributes associated with the route suggests that the signaling protocol traffic will not traverse through certain ITADs. This can be ensured by checking the RoutedPath attribute for the prohibited ITADs as soon as the routing update is received on the ETRIP session.

This ability to filter routes based on attributes (for both incoming and outgoing updates) can be very helpful in planning and optimizing network designs for capacity planning and fault-tolerance. It can also come handy in enforcing settlement rules between administrative domains.

Security in TRIP

Since TRIP LSs peer with each other to exchange routing information that is critical to the correct and efficient operation of an ITAD, it is important that TRIP safeguard itself from any unauthorized peering or "sniffing" of the information being exchanged. TRIP suggests the use of IPSec to secure information. IPSec is a standardized security mechanism available for IP networks.

TRIP also takes care of protecting TRIP Routes from being tempered by an unauthorized entity. In many situations, an LS receives a route, which has been originated by remote LS that is not a direct peer of the receiving LS. In addition, attributes may have been inserted or altered along the advertisement path. The receiving LSs may wish to authenticate the route by verifying both the originator of an attribute and fact that the contents of the attribute have not been altered by other intermediate LSs. The Authentication attribute carries a list of signatures so that a receiving LS may validate particular attributes.

TRIP in an H.323 Network

In this section, we address how TRIP can be incorporated into the H.323 world to offer a call routing solution. We first look at how the H.323 network would look like and then demonstrate using sample call flows, how TRIP can be used to route a call.

An Internet Telephony Administrative Domain (ITAD), as it applies to H.323, can be thought of as constituting one or more H.323 zones. An ITAD encompassing an H.323 network can have one or more TRIP

Location Servers (LS), H.323 Gateways, one or more Gatekeepers and Annex G Border Element(s). These Location Servers could exist as stand-alone entities in the network or could be co-located with one of the H.323 network entities (for example, with a Gatekeeper)

When a call is to be placed, it is possible for any H.323 entity to query the TRIP LS to fetch the best route for that telephony prefix. The protocol to query a TRIP speaker is unspecified. A RAS-like protocol could be considered for this purpose. Also, the H.323 entity to interface with the TRIP LS should be determined based on the design of the H.323 network.

In response to a query, a TRIP LS could return a route specifying one of the following H.323 signaling protocols to be used to contact the Signaling Server on the next hop -- H.323-Q.931, H.323-RAS or H.323-AnnexG. This means that to progress with the call set up for that telephony prefix, the querying entity would be required to establish a session with the Signaling Server provided using the specified signaling protocol.

The rest of the document discusses some call scenarios demonstrating the use of TRIP in routing a call in an H.323 Network.

Call Flow with H.323-Q.931 in a TRIP route

Figure 2 shows ITADs ITAD1 and ITAD2, each having a TRIP Location Server. The gateways in each ITAD register the telephony prefixes that they can terminate with the corresponding LS. In addition, the gateways are setup to query the TRIP LS when a call is to be placed. We have a sample call flow demonstrated below, where a call from the PSTN is handled by querying TRIP by the gateway. The following are the steps taken to route a call in this scenario.

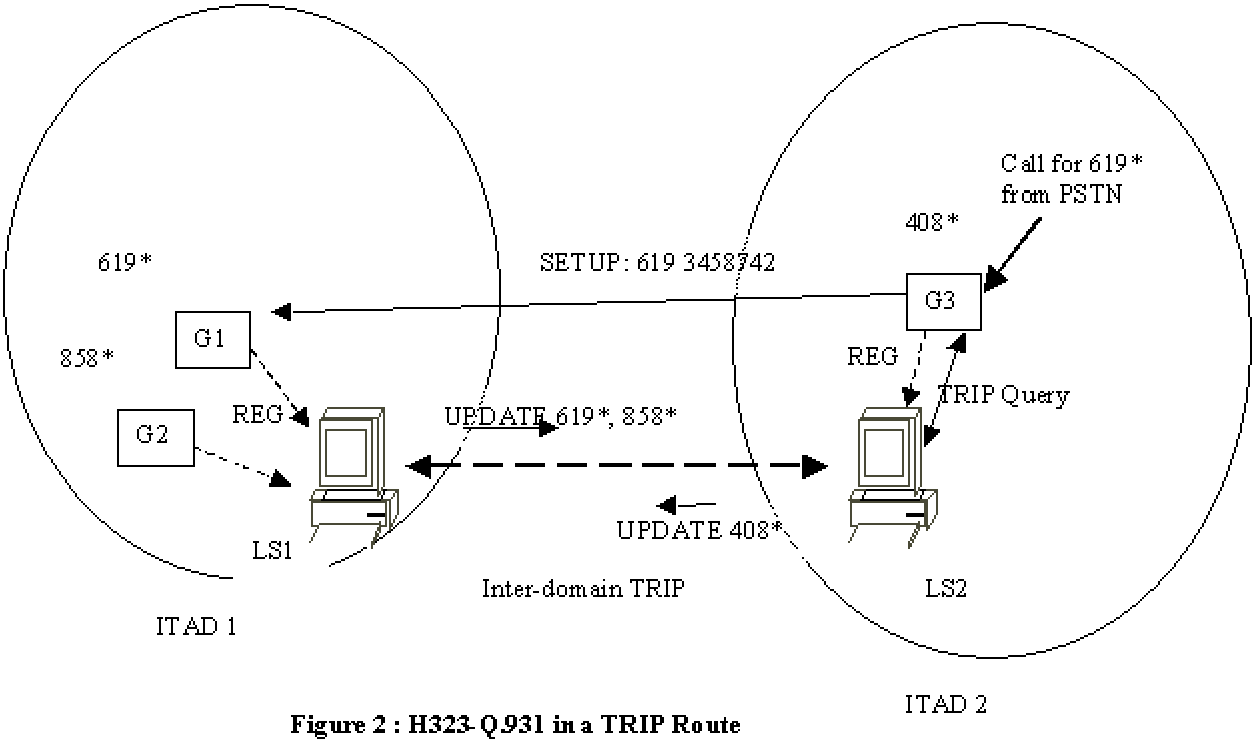

Figure 2 : H323-Q.931 in a TRIP Route

1. Gateway G3 in ITAD2 receives a call for prefix 619* from the PSTN
2. G3 queries LS2
3. LS2 returns a route {NextHop:G1, Signaling:H323-Q931}
4. G3 sends a Q.931 SETUP message to Gateway G1

Call Flow with H323-RAS in a TRIP route

In this case, we consider a topology similar to the previous case, but with the inclusion of Gatekeepers (Figure 3). The network consists of two H.323 zones, with ITAD1 encompassing one zone and ITAD2 encompassing the other. Each zone has a gatekeeper that interacts with the Location Server in its ITAD. The gateways are setup to consult the gatekeeper for routing a call. In this sample call flow, we demonstrate use of a TRIP route that specifies use of H323-RAS to reach the next hop signaling point for a call received from the PSTN.

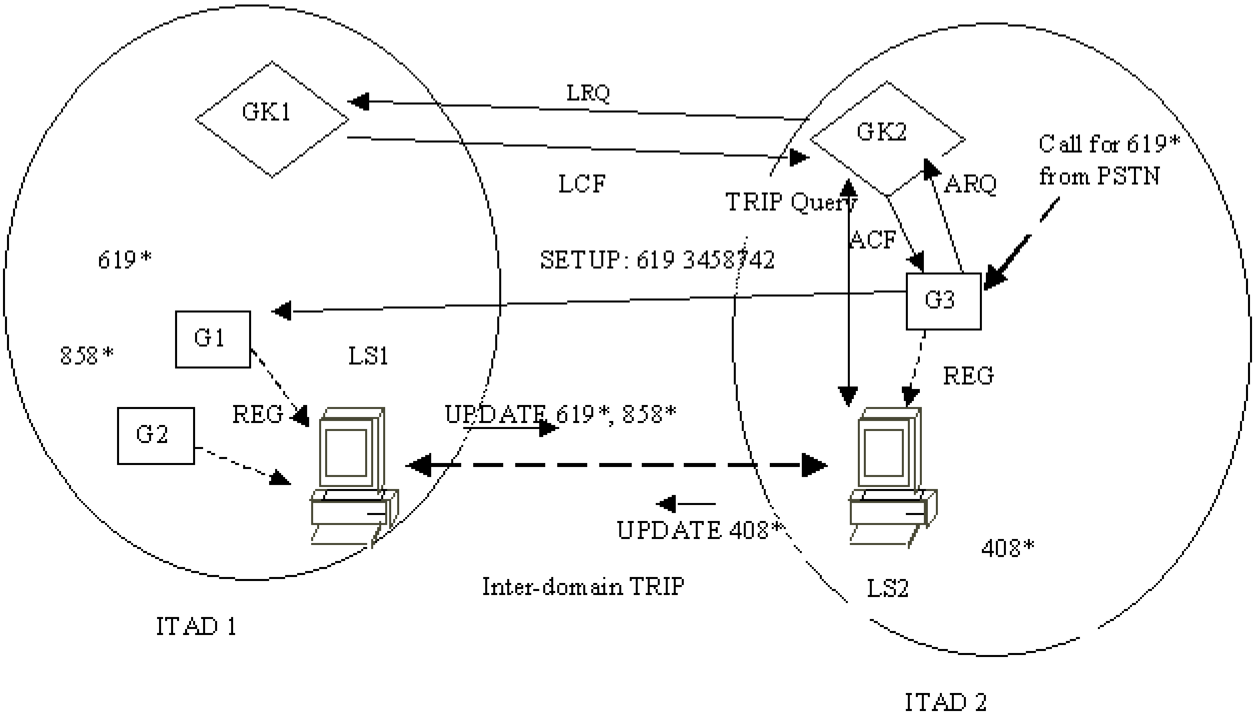

Figure 3 : H.323-RAS in a TRIP Route

1. Gateway G3 in ITAD2 receives a call for prefix 619* from the PSTN
2. G3 sends an ARQ to GK2
3. GK2 sends a TRIP Query to LS2
4. LS2 returns a route {NextHop:GK1, Signaling:H323-RAS}
5. GK2 sends a RAS LRQ to GK1
6. GK1 responds with LCF with IP address of Gateway G1
7. GK2 sends an ACF to Gateway G3 with IP address of Gateway G1
8. G3 sends a Q.931 SETUP message to Gateway G1

The above call flows show a stand-alone TRIP LS in each ITAD. However, it is possible to have a TRIP LS co-locate with an H.323 gatekeeper.

Implementing IP Routing

Archived content. No warranty is made as to technical accuracy. Content may contain URLs that were valid when originally published, but now link to sites or pages that no longer exist.

By Todd Lammle,with Monica Lammle and James Chellis

Chapter 3 from *MCSE: TCP/IP for NT Server 4 Study Guide*, published by Sybex, Inc.

With TCP/IP basics covered and conquered in Chapters 1 and 2, our focus is going to both sharpen and shift. Our attention will now be concentrated on Microsoft-specific issues and intricacies. We will begin Chapter 3 with a crisp discussion on IP routing.

On This Page

Objectives

Our topic scope will encompass several important objectives. By the end of this chapter, you should be able to do the following:

- Explain the difference between static and dynamic IP routing.
- Explain the host configuration requirements to communicate with a static or dynamic IP router.
- Configure a computer running Windows NT to function as an IP router.
- Build a static routing table.
- Use the TRACERT utility to isolate route or network link problems.

Tip All readers—even the wizards and gurus in the audience—shouldn't skip class today. Like I said, we're moving into Microsoft Land now, and this chapter is fundamental because it deals with IP routing as spoken in MSNT—an important Microsoft dialect.

Top Of Page

What Is IP Routing?

IP routing is the process of sending data from a host on one network to a *remote host* on another network through a *router*, or routers. A router is either a specifically assigned computer or a workstation that's been configured to perform routing tasks. In IP terminology, routers are referred to as *gateways*. Gateways are basically TCP/IP hosts that have been rigged with two or more network connection *interfaces*. Outfitted in this manner, they're known as *multihomedhosts,* which we'll discuss more thoroughly later in the chapter.

The path that a router uses to deliver a packet is defined in its *routing table*. A routing table contains the IP addresses of router interfaces that connect to the other networks the router can communicate with. The routing table is consulted for a path to the network that is indicated by the packet's destination address. If a path isn't found, the packet is sent to the router's *default gateway* address—if one is configured. By default, a router can send packets to any network for which it has a configured interface. When one host attempts communication with another host on a different network, IP uses the host's default gateway address to deliver the packet to the corresponding router. When a route is found, the packet is sent to the proper network, then onward to the destination host. If a route is not found, an error message is sent to the source host.

The IP Routing Process

The IP routing process is fairly direct when a datagram's destination is located on a neighboring network. In this kind of situation, a router would follow a simple procedure, as shown in Figure 3.1.

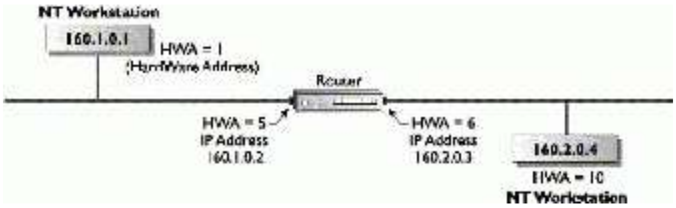

Figure 3.1: Simple routing

First, when a workstation wants to send a packet to a destination host, in this instance 160.1.0.1 transmitting to 160.2.0.4, host 160.1.0.1 checks the destination IP address. If it determines the address isn't on the local network, it must then be routed. Next, 160.1.0.1 calls on ARP to obtain the hardware address of its default gateway. The IP address of the default gateway is configured in machine 160.1.0.1's internal configuration, but 160.1.0.1 still needs to find the hardware address of the default gateway, and sends out an ARP request to get it. IP then proceeds to address the packet with the newly obtained destination hardware address of its default router. The information utilized for addressing the packet includes:

- Source hardware address 1
- Source IP address 160.1.0.1
- Destination hardware address 5
- Destination IP address 160.2.0.4

IP, on the receiving router with the hardware address of 5, establishes that it is not the final, intended recipient by inspecting the packet's destination IP address, which indicates it must be forwarded to network 160.2. Then, IP uses ARP to determine the hardware address for 160.2.0.4. The router then puts the newly identified hardware address into it's ARP cache for easy reference the next time it's called upon to route a packet to that destination.

This accomplished, the router sends the packet out to network 160.2 with a header that includes:

- Source hardware address 5
- Source IP address 160.1.0.1
- Destination hardware address 10
- Destination IP address 160.2.0.4

As the packet travels along network 160.2, it looks for hardware address 10, with the IP address of 160.2.0.4. When an NIC card recognizes its hardware address, it grabs the packet.

It's important to note here that the source IP address is that of the host that created the packet originally, but that the hardware address is now that of the router's connection interface to network 160.1. It's also significant that although both source and destination software IP addresses remain constant, both source and destination hardware addresses necessarily change at each hop the packet makes.

Sounds simple right? Well, it is in a situation like the one we just presented. However, those of you who possess some firsthand experience with this sort of thing may now be finding yourselves just a little distracted with thoughts of a genuinely sarcastic variety. Before turning your nose up and slamming this book shut, let it be known that we, too, are fully aware that this isn't a perfect world, and that if things were that straightforward, there wouldn't be a market for books about them! On the other hand, to those readers becoming uncomfortable with the now present implications of potential chaos, we say, relax, make some tea, sit down, and read on.

Start by considering this foul and ugly possibility: What if the destination network is in the dark because it's not directly attached to a router on the delivery path for that nice little datagram? Things come a tumblin' down, that's what! Remember hearing somewhere that there's a reason for everything? Well, we're not sure about that, but the heinous, confusion-producing scenario we just posited is one most excellent reason for the existence of routing tables. With a handy-dandy routing table, the fog clears, clouds part, and destinations sing! Routers, and those dependent on them, again become happy, efficient things. Routing tables are maintained on IP routers. IP consults these to determine where the mystery network is, so that it can send its mystery packet there. Some internetworks are very complex. If this is the case, routing tables should designate all available routes to a destination network, as well as provide an estimate advising the efficiency of each potential route. Routing tables maintain entries of where networks are located, not hosts.

Dynamic vs. Static IP Routing

There are two breeds of routing tables. There are static tables, and there are dynamic tables. Static types are laboriously maintained by a network manager, while the dynamic variety is sustained automatically by a routing protocol. Here's a list spotlighting some specific routing table characteristics:

Dynamic Routing	Static Routing
Function of inter-routing protocols	Function of IP
Routers share routing information automatically	Routers do not share routing information
Routing tables are built dynamically	Routing tables are built manually
Requires a routing protocol, such as RIP or OSPF	Microsoft supports multihomed systems as routers
Microsoft supports RIP for IP and IPX/SPX	

Windows NT Server 4.0 provides the ability to function as an IP router using both static and dynamic routing. A Windows NT-based computer can be configured with multiple network adapters and can route between them. This type of system, which is ideal for small, private internetworks, is referred to as a multihomed computer.

Windows NT Server 4.0 has the ability to function as a Routing Information Protocol (RIP) router that supports dynamic management of Internet Protocol (IP) routing tables. RIP eliminates the need to establish static IP routing tables.

Dynamic IP Routing

On large internetworks, dynamic routing is typically employed. This is because manually maintaining a static routing table would be overwhelmingly tedious, if not impossible. With dynamic routing, minimal configuration is required by a network administrator. Figure 3.2 shows an example of dynamic routing.

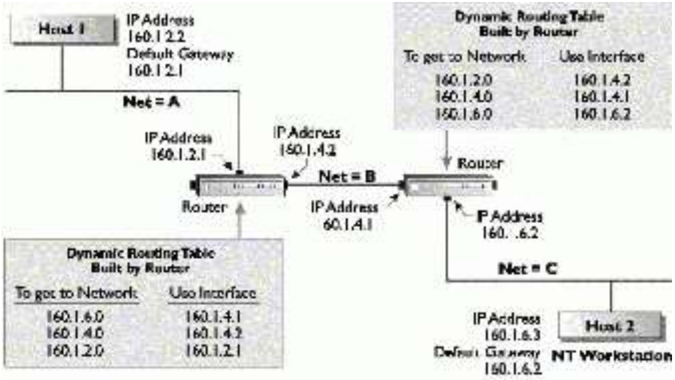

Figure 3.2: An example of dynamic routing

For a host to communicate with other hosts on the internetwork, its default gateway address must be configured to match the IP address of the local router's interface.

In Figure 3.2, Host 1 requires a default gateway address in order to be able to send packets to any network other than Network A. Host 1's default gateway address is configured for the router port attached to Network A. If Host 1 sends a packet that's not destined for the local network, it will be sent to the default gateway address. If no gateway is defined, the packet will be discarded. Host 2 works the same way; however, it's default gateway is the router port attached to Network C. When the router receives a packet either from Host 1 or 2, it will observe the destination's IP address and forward it according to the information in its routing table, which is built and maintained through inter-routing protocols.

Dynamic routing is a function of inter-routing, network gossip protocols such as the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). These routing protocols periodically exchange routes to known networks among dynamic routers. If a given route changes, they automatically update the router's routing table and inform other routers on the internetwork of the change.

Routing Information Protocol (RIP)

RIP is a type of protocol known as a *distance vector routing protocol.* RIP is used to discover the cost of a given route in terms of hops, and store that information in the routing table, which IP uses in selecting the most efficient, low-cost route to a destination. It works by watching for routing table broadcasts by other routers and updating its own routing table in the event a change occurs. RIP routing tables provide, at a minimum, the following information:

- IP destination address
- A metric (numbered from 1 to 15) indicative of the total cost in hops of a certain route to a destination
- The IP address of the router a datagram would reach next on the path to its destination
- A marker signaling recent changes to a route
- Timers

Some drawbacks to RIP include a problem known as "counting to infinity," as illustrated in Figure 3.3. In certain internetwork configurations, an endless loop between routers can occur if one of the networks becomes unavailable. RIP keeps counting hops each time the broadcast reaches a router, in hopes of finding a new route to the formerly available network. To prevent this, a hop-limit count between 1 and 15 is configured to represent infinity, which necessarily imposes size restrictions on networks. RIP can't be utilized on a network with an area consisting of more than 15 hops. In Figure 3.3, Network 6's location was lost between Routers B and D. Router B then looks for a new route to Network 6. Router B already knows that Router C can get to Network 6 with four hops because Router C advertised this information in a broadcast, and all routers save this broadcasted information in their routing tables. Since Router B is looking for a new route to Network 6, Router B references its routing table and finds that Router C can reach Network 6 in four hops. Router B determines it can reach Network 6 in five hops because Router C can make it in four hops. This is because Router B must add an extra hop for itself—four from Router C plus one for Router B. Router B then broadcasts the new route information back out onto the network. Router C receives this information, and enters into its route table that Network 6 is now six hops away—five from Router B, plus one for itself. This process continues until the 15-hop limit is reached. At this point, the route to Network 6 is finally dubbed an unreachable destination, and all related route information regarding it is removed from both Router B's and Router C's routing tables.

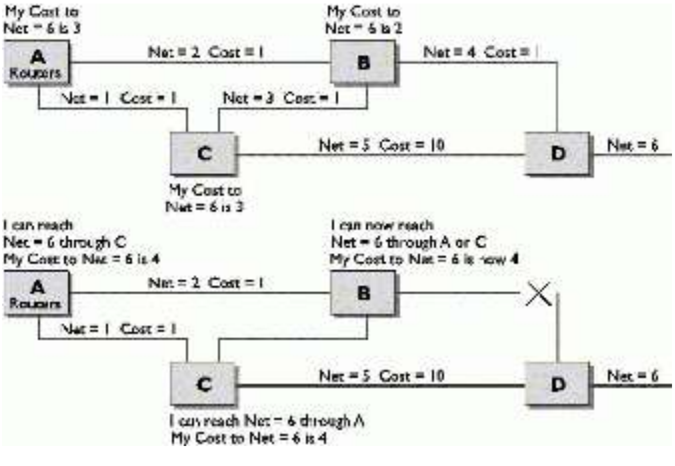

Figure 3.3: "Counting to infinity"

Another problem with large internetworks centers around the fact that RIP routers broadcast routing table advertisements every 30 seconds. On today's gargantuan networks, populated with an abundance of routers, momentous amounts of bandwidth can get gobbled up simply accommodating all the RIP response packet noise.

Open Shortest Path First (OSPF)

Because of these potentially network-hostile characteristics, OSPF is quickly gaining popularity within the Internet community. OSPF is based on *link-state algorithms*, and is therefore known as a *link-state routing protocol*. It's deployed within an *autonomous system*, which is a group of routers that share a certain routing protocol. When that protocol happens to be OSPF, each router retains its own database describing the topology of the autonomous system on which it's located. This kind of system is much more flexible, and has the following advantages as well:

- Network administrators can assign costs to a particular link.
- The total cost for a given path doesn't necessarily have to have a limit.
- Its upper metric limit being 65,535, it has the ability to accommodate vast networks.
- Each node creates a link-state database tree representing the network and places itself as that tree's root, where it can choose the most direct path to a given destination.
- Related to the above benefit, in the event that more than one route of equal cost exists, OSPF routers can balance the load of network traffic between all available and equally cost-effective routes.
- Link-state routing advertisements are broadcasted much less often—only when a change is detected, thereby reducing network overhead.
- Link-state routing update packets can efficiently carry information for more than one router.
- This type of packet is only sent to *adjacencies,* or neighboring routers selected to swap routing information—a "tell a friend" arrangement that again contributes to network efficiency.

Static IP Routing

Static routing is a function of IP. Static routers require that routing tables be built and updated manually. If a route changes, static routers are secretive and do not share this information to inform each other of the event. They're also very cliquey, and do not exchange routes with dynamic routers.

Windows NT can function as an IP router using static routing. NT network administrators must maintain their tables or acquire a commercial router. A Windows NT-based computer can be configured with multiple network adapters and routes between them. This type of system, which is ideal for small, private internetworks, is referred to as a *multihomed computer.*

Routing tables inventory known networks and the IP addresses used to access them. Windows NT static routing tables are maintained by a route utility, and are comprised of five columns of data, reading left to right. In the list below, the first entry represents the left-most column, the second represents the next one, and so on.

Network Address: A roster of addresses for known networks. Included here is an entry both for the local network (0.0.0.0) and for broadcasts (255.255.255.255).

Netmask: This column lists all subnet masks in use for each network.

Gateway Address: This is a list of the IP addresses employed as the primary datagram receivers for each network.

Interface: Each network card installed in a computer is assigned an interface number.

Metric: This is a list providing an estimate of the number of hops the route would cost. A hop is each pass a datagram makes through a router.

Here are the key things to remember about static routers:

- A static router can only communicate with networks with which it has a configured interface.
- A Windows NT computer can be configured as a multihomed computer.
- A static route can be configured as either a default gateway address or an entry in a routing table.
- A static router, such as a Windows NT multihomed computer, can only communicate with networks to which it has a configured interface. This limits communications to local networks.

Figure 3.4 illustrates static routing.

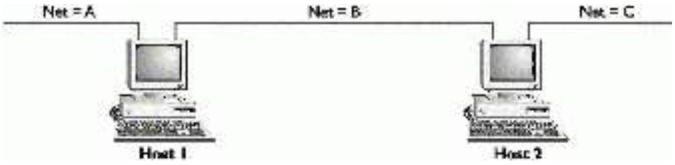

Figure 3.4: Static routers

As shown in Figure 3.4, Host 1 has local connections to Networks A and B. This means that hosts on Network A can communicate with hosts on Network B, and vice versa, because Host 1 knows about both networks and will pass packets destined for either one. Hosts on Network A will not be able to communicate with hosts on Network C.

Host 2 has local connections to Networks B and C and will be able to pass packets destined for either network. However, Network C will not be able to send packets to Network A.

Note: After taking all this in, you may have been left with the impression that dynamic routing is the method of choice for everyone's routing needs. While that's certainly true when the network in question is large and complex, providing a multiplicity of paths to destinations and/or growing rapidly, static routing is wonderfully suited for small to moderately sized networks that rarely change. An important consideration is, as is so often the case, cost. All that fabulous intelligence and flexibility, and all those bells and whistles that dynamic routers offer cost a lot—up to around $50k or more apiece! They're one of the most expensive pieces of equipment one can hook up to a network! Windows NT comes out of the box equipped with static routing built right in—in other words...it's free! It's also free of charge in terms of overhead costs on your network, and it creates the environment for a much closer, more involved relationship between you and your beloved network.

Top Of Page

IP Routing Applied

Now that you have a clear picture of exactly what IP routing is, and what it involves, you're ready to learn how it's done. In this section, we'll give you the skinny on configuration and integration issues, and the procedures required for implementation.

Using the Default Gateway Address on a Static Router

Gateways are most often dedicated computers, or routers. The *default gateway* is like a network mediator with connections. It's the node on the local network that knows the network IDs of other networks linked to the greater internetwork. Since it has access to this privileged information, when a given workstation sends out some data that reaches the default gateway, it can forward it along to other gateways as required to reach its proper destination.

One method of configuring a static route without manually adding routes to a routing table is to configure each multihomed computer's default gateway address as the local interface to the other multihomed computer on the common network. It's a type of circular reasoning for computers.

A multihomed computer (a computer with more than one NIC card) can send IP packets to destinations other than those they are locally attached to by setting the internal configuration of the default gateway to the other multihomed computer's network interface. For example, in Figure 3.4, Host 2 would set its default gateway to the network interface on Host 1. Network C then would be able to pass packets to Network A. Host 1 would set its default gateway to the network interface on Host 2, enabling Network A to communicate with Network C.

Whenever Host 1 receives a packet destined for a host on Network C, it'll check its local routing table. If it doesn't find a route to Network C, it forwards the packet to its default gateway, which is a local interface on Host 2. Host 2 will then route the packet to the appropriate interface for delivery on Network C.

Note: Using the default gateway address as a static route only works well with two routers. If more than two are used, you must manually add an entry in the routing table.

Using Additional Default Gateways

Although more than one default gateway can be configured, only the first one will be used for routing purposes. The others will be used only as backups should the primary one become unavailable for some reason. This means you can't use multiple gateways to nab more network bandwidth. However, better fault tolerance is still a definite plus. As I'm sure you are well aware, this backup stuff is by no means unimportant.

Let's say Router A in Figure 3.5 goes on the blink and is out of commission when the client boots up. The client wants to connect to the server, but the default gateway that the client is defined for is currently unavailable, so it's out of luck. The client has no other default gateway defined, so the client will not be redirected to Router C to connect to the server. However, if the client was to define a second default gateway for Router C, the client would then be directed on towards the server.

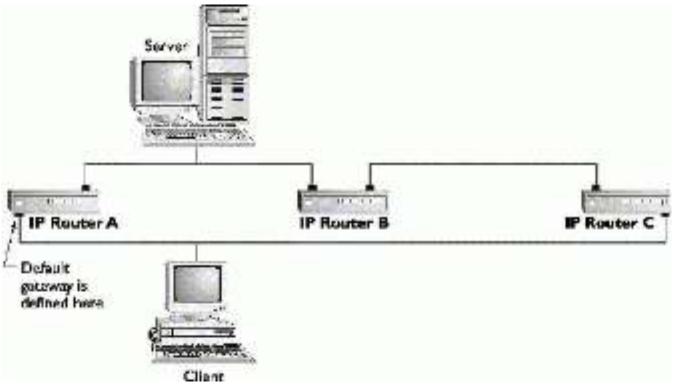

Figure 3.5: A hypothetical routing dilemma

The client is going to have to sit on the other side of the abyss dreaming of multiple default gateways, wishing he or she had a more thorough network administrator who took these kind of precautions. Would that be you? If so, Exercise 3.2 shows you how to create additional default gateways using the Advanced TCP/IP configuration.

When one configures one's Windows NT system in this way, retransmission problems at the TCP layer will cause the IP routing software to try the routers listed in the Additional Gateways value. This is again a great backup plan. Why? Well…what if Client and Server were in the middle of an established session, and that troublesome Router A went down again? TCP would send a message to IP to try one of the additional routers in the registry. IP would then try Router C, use a double hop route, and continue exchanging data at no cost to the session. When Router A again breathes a breath of life, the inter-routing protocol will force Router C to redirect the session's traffic back through the more optimal, one-hop route provided by the now living Router A.

Note: In order for any host to be able to communicate with other hosts located somewhere out there on the internetwork, its default gateway address positively must be configured to match the IP address of the router's interface on the local segment.

Default Routing Table Entries

Windows NT 4.0, NT 3.51, and Windows for Workgroups 3.11 routing tables maintain the default entries shown in Table 3.1.

Table 3.1 Default routing tables for Windows NT 4.0, 3.51, and Windows for Workgroups 3.11

Address Examples	Description
0.0.0.0	The address used as a default route for any network not specified in the route table
Subnet broadcast	The address used for broadcasting on the local subnet
Network broadcast	The address used for broadcasting throughout the internetwork
Local loopback 127.0.0.1	The address used for testing IP configurations and connections
Local network	The address used to direct packets to hosts on the local network
Local host	The address of the local computer; references the local loopback address

Adding Static Entries

The route command is used, as detailed in Table 3.2, to add static entries to a routing table.

Table 3.2 Commands Used for Adding Static Entries

To add or modify a static route	Function
route add [network address] mask [gateway address]	Adds a route
route delete [network address] [gateway address]	Deletes a route
route change [network address] [gateway address]	Modifies a route
route print [network address] [gateway address]	Prints a route
route -s [gateway address]	Adds a route to a smart gateway
route -f	Clears all routes

It works like this: To add a static route, enabling communications between a host on network 160.1.89.0 from a host on network 160.1.66.0, you would run the following command:

```
route add 160.1.24.0 mask 255.255.255.0 160.1.16.1
```

Note: Static routes are stored in memory unless the -p parameter is used. No, it doesn't stand for permanent; it stands for persistent. Persistent routes are stored in the registry. If you restart your NT Server or workstation, you will have to re-create all nonpersistent routes if you're using a static routing table.

If your internetwork has more than two routers, at least one of which is a static router, you'll need to configure static routing table entries to all known networks on a table at each multihomed computer, as shown in Figure 3.6.

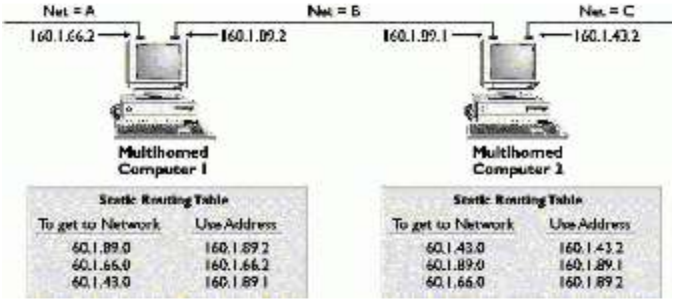

Static Routing Table

To get to Network	Use Address
60.1.89.0	160.1.89.2
60.1.66.0	160.1.66.2
60.1.43.0	160.1.89.1

Static Routing Table

To get to Network	Use Address
60.1.43.0	160.1.43.2
60.1.89.0	160.1.89.1
60.1.66.0	160.1.89.2

Figure 3.6: Proper routing table configuration

A static routing table for Network C, IP address 160.1.43.0, is created on multihomed Host 1, directing the router to send the packets to interface 160.1.89.1. This will permit packets to move from Network A to Network C.

Also, a static route is configured on Host 2 to allow Network C to send packets to Network A. This is accomplished by referencing the network address 160.1.66.0, directing the router to send packets on to IP address 160.1.89.2, which can directly deliver packets destined for Network A.

The static routing table is always checked before a packet is routed. If there is no static route to a particular host, the packet is sent to the configured default gateway. For a host to communicate with other hosts on the internetwork, its default gateway address must be configured to match the IP address of the local router's interface.

Integrating Static and Dynamic Routing

As mentioned above, breeds of routers stick to their own and do not speak with those of a different feather. Static routers do not trade routing information with dynamic routers unless they're forced to. As they say, where there's a will, there's a way! To route from a static router through a dynamic router, such as a RIP or OSPF-enabled IP router, one must first add a static route to the routing tables located on both the static and dynamic routers, as shown in Figure 3.7.

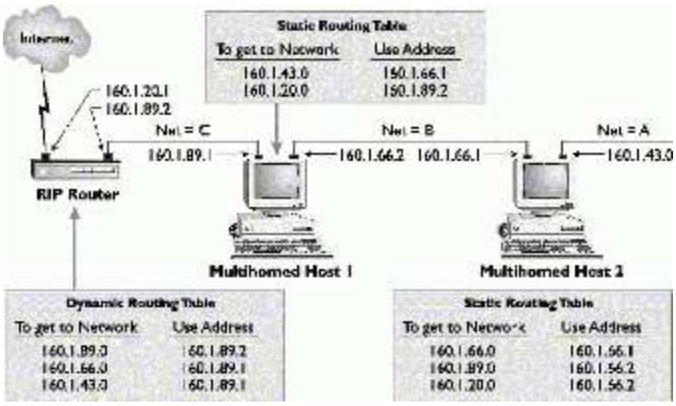

Static Routing Table

To get to Network	Use Address
160.1.43.0	150.1.66.1
160.1.20.0	150.1.89.2

Dynamic Routing Table

To get to Network	Use Address
160.1.89.0	160.1.89.2
160.1.66.0	160.1.89.1
160.1.43.0	160.1.89.1

Static Routing Table

To get to Network	Use Address
160.1.66.0	160.1.56.1
160.1.89.0	160.1.56.2
160.1.20.0	160.1.56.2

Figure 3.7: Static and dynamic router integration

In order to route packets from Network A to the Internet, Host 2 requires that a route be added to its routing table. This addition includes the IP address of the closest interface (in this case, 160.1.66.2) that can access the dedicated IP router to the Internet.

To route packets from Networks B and C to the Internet, a static entry must be added to Computer B's routing table. This entry includes the IP address of the nearest interface (160.1.89.2) on the dedicated IP router to the Internet.

To allow computers on the Internet to communicate with hosts on Networks 1 and 2, one must statically configure the dedicated IP router with the IP address of the interface to Host 1. Host 1 then acts as a gateway to other subnets.

Note: It's important to mention here that some implementations of RIP do not propagate static routing tables. Should you find yourself beset with this dilemma, you'd have to statically configure the remote routers in the Internet cloud to achieve your routing goals. Also important to note here is that the exact method for configuring a static route on a RIP router varies with each kind of router. It is therefore very wise to refer to your particular router's vendor documentation for more information.

In Exercise 3.1, you'll view and configure the routing table.

Exercise 3.1: Viewing and Configuring the Routing Table

1. From a command prompt, type **route print**, and then press Enter to view the route table.

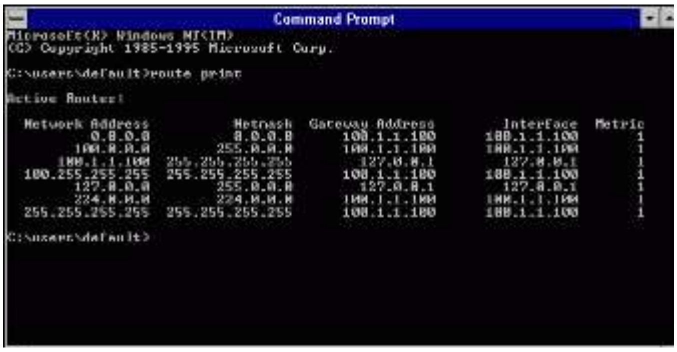

```
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1995 Microsoft Corp.

C:\users\default>route print

Active Routes:

  Network Address          Netmask  Gateway Address        Interface  Metric
          0.0.0.0          0.0.0.0      100.1.1.100      100.1.1.100       1
        100.0.0.0        255.0.0.0      100.1.1.100      100.1.1.100       1
      100.1.1.100  255.255.255.255        127.0.0.1        127.0.0.1       1
  100.255.255.255  255.255.255.255      100.1.1.100      100.1.1.100       1
        127.0.0.0        255.0.0.0        127.0.0.1        127.0.0.1       1
        224.0.0.0        224.0.0.0      100.1.1.100      100.1.1.100       1
  255.255.255.255  255.255.255.255      100.1.1.100      100.1.1.100       1

C:\users\default>
```

2. Under Gateway Address, you should find your default gateway address—the router interface address attached to your local network interface.
3. Remove your default gateway address from your computer. This will prevent any packets being sent to the default gateway for routing and require all routing to be done from existing route entries.
4. Access the Microsoft TCP/IP Properties dialog box.
5. Delete the Default Gateway address.
6. Click OK in the Microsoft TCP/IP Properties box.
7. Click OK.
8. Switch to a command prompt, and use the route print command. The default gateway should not be listed.

Next, attempt communication with both local and remote hosts.

1. Ping the IP address of a local host. This should be successful.
2. Ping the IP address of a host on a remote address. This should fail, with a "Destination host unreachable" error.
3. Add a route entry to your computer.
4. Type the following command: **route add *remote_network_id* mask *subnet_mask* *your_default_gateway*** . (For example, if your workstation is on network 150.150.28.0 and the remote network is 150.150.40.0, the command would look like this: **route add 150.150.40.0 mask**

255.255.252.0 150.150.28.200. This means that to get to network 150.150.40.0, you would use address 150.150.28.200, which should be the default gateway.)

5. View the route table.
6. Ping a remote host. This should be successful.
7. Restore your default gateway.
8. Test communication by pinging your default gateway.

Implementing a Windows NT Router

Windows NT originally shipped with the ability to act as a static IP router. The static IP router was enabled by creating a multihomed system and enabling routing either through the registry or through a checkbox in the Advanced TCP/IP Configuration dialog box.

Static routing can work well for small networks and remote sites, but for large internetworks, the overhead of manually maintaining routing tables is significant. By enabling the RIP for IP routing protocol, Windows NT Server 4.0 can be a dynamic IP router. Windows NT 4.0 RIP for IP eliminates the manual configuration of routing tables. RIP for IP is suitable for medium-size internetworks, but it is not suitable for large IP internetworks because of the significant amount of broadcast traffic it generates.

To implement a Windows NT Router:

1. Install multiple adapter cards and appropriate drivers, or configure multiple IP addresses on a single card.
2. Configure the adapter card(s) with a valid IP address and subnet mask.
3. On the Routing tab of the Microsoft TCP/IP Properties dialog box, select the Enable IP Forwarding checkbox.

   Depending on which version of Windows NT you are running:

   - On the Services tab of the Control Panel Network tool, add the RIP for Internet Protocol service.
   - Add static routes to the static routers routing table for all networks to which the computer has no configured interface.

The TRACERT Utility

The *TRACERT utility* is essentially a verification tool. It's used to substantiate the route that's been taken to a destination host. TRACERT is also highly useful in isolating routers and identifying WAN Links that are not functioning and/or are operating too slowly. Here's the relevant command syntax for deploying TRACERT:

```
tracert 160.1.89.100
```

where 160.1.89.100 is the remote computer.

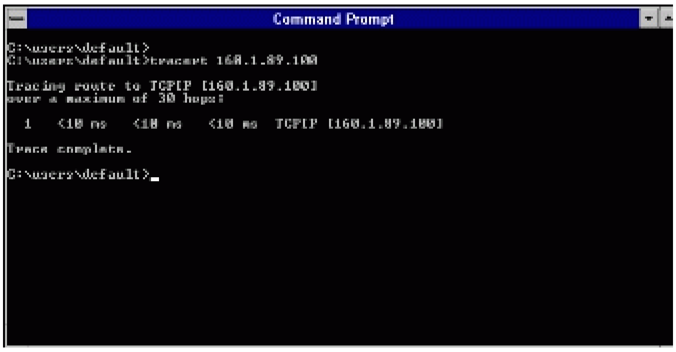

```
C:\users\default>
C:\users\default>tracert 160.1.89.100

Tracing route to TCPIP [160.1.89.100]
over a maximum of 30 hops:

  1    <10 ms    <10 ms    <10 ms  TCPIP [160.1.89.100]

Trace complete.

C:\users\default>_
```

Below is an example of the output that would result from entering the command if the computer was on the network:

```
Tracing route to 160.1.89.100 over 2 hops1 <10 ms <10 ms <10 ms¬
160.1.89.22
<10 ms <10 ms <10 ms 160.1.89.100
```

TRACERT can aid in determining if a router has failed by the degree of success the command enjoys. For example, if the command is unsuccessful, it is possible to assess router or WAN link problems and identify the point at which routing failed. The response time for the command is returned in the output. The information contained in the output can be readily compared to that recorded for another route to the same destination. This greatly facilitates identifying a slow or ineffective router or WAN link.

For example, the command tracert 160.1.66.11 would display the path taken from the local host to the destination host: 160.1.66.11. The output from the preceding command would confirm that the router address 160.1.66.1 was the route taken from the local host to the destination host. Here's what that output would look like:

```
Tracing route to 160.1.66.11 over a maximum of 30 hops1 <10 ms¬ <10 ms <10
ms
160.1.66.12 <10 ms <10 ms <10 ms 160.1.66.11
```

Let's go to the Internet and do a tracert to the White House. This is the output we receive:

```
C:\WINDOWS>tracert www.whitehouse.gov
Tracing route to www.whitehouse.gov [198.137.240.92]
over a maximum of 30 hops:
1 191 ms 136 ms 138 ms 38.1.1.1
2 142 ms 132 ms 137 ms 38.17.3.1
3 157 ms 151 ms 152 ms 38.1.42.3
4 152 ms 149 ms 147 ms 38.1.42.3
5 206 ms 285 ms 338 ms se.sc.psi.net [38.1.3.5]
6 286 ms 290 ms 265 ms rc5.southeast.us.psi.net [38.1.25.5]
7 278 ms 317 ms 268 ms ip2.ci3.herndon.va.us.psi.net [38.25.11.2]
8 268 ms 264 ms 313 ms 198.137.240.33
9 288 ms 279 ms 346 ms www.whitehouse.gov [198.137.240.92]
Trace complete.
C:\WINDOWS>
```

IPv6 (IPng)

This protocol used to be referred to as "IP Next Generation" or "IPng." IP version 6 (IPv6) is a new version of the Internet Protocol designed as a successor to IP version 4 (Ipv4, RFC-791). The current header in IPv4 hasn't been changed or upgraded since the 1970s! The initial design, of course, failed to anticipate the growth of the Internet, and the eventual exhaustion of the Ipv4 address space. Ipv6 is an entirely new packet structure, which is incompatible with IPv4 systems.

Expanded Addressing Capabilities

IPv6 has a128-bit source and destination IP addresses. With approximately five billion people in the world using 128-bit addresses, there are 2128 addresses, or almost 296 addresses per person! An IPv6 valid IP address will look something like this:

```
3F3A:AE67:F240:56C4:3409:AE52:220E:3112
```

IPv6 uses 16 octets; when written, it is divided into eight octet pairs, separated by colons and represented in hexadecimal format.

Header Format Simplification

The IPv6 headers are designed to keep the IP header overhead to a minimum by moving nonessential fields and option fields to extension headers after the IP header. Anything not included in the base IPv6 header can be added through IP extension headers placed after the base IPv6 header.

Improved Support for Extensions and Options

IPv6 can easily be extended for unforeseen features by adding extension headers and option fields after the IPv6 base header. Support for new hardware or application technologies is built in.

Flow Labeling Capability

A new field in the IPv6 header allows the pre-allocation of network resources along a path, so time-dependent services such as voice and video are guaranteed a requested bandwidth with a fixed delay.

Note: Not to worry, they don't ask you to subnet IPv6 on the NT4.0 TCP/IP test.

Top Of Page

Supernetting

The Internet is running out of IP addresses. To prevent the complete depletion of network IDs, the Internet authorities have come up with a scheme called *supernetting.* The main difference between subnetting and supernetting is that supernetting borrows bits from the network ID and masks them as the host ID for more efficient routing. For example, rather then allocate a Class B network ID to a company that has 1800 hosts, the InterNIC allocates a range of eight Class C network IDs. Each Class C network ID gives 254 available hosts for a total of 2,032 host IDs.

However, this means that the routers on the Internet now must have an additional eight entries in their routing tables to route IP packets to the company. To prevent this problem, a technique called *Classless Inter-Domain Routing* (*CIDR*) is used to collapse the eight entries used in the above example to a single entry corresponding to all of the Class C network IDs used by that company.

Let's take a look at how this would work with eight Class C networks IDs starting with the network ID 220.78.168.0 through the network ID 220.78.175.0. In the routing table the entry then would be:

Network ID	Subnet Mask	Subnet Mask Binary
220.78.168.0	255.255.248.0	11111111.11111111.11111000.00000000

A typical Class C network would have a subnet mask of 255.255.255.0. Only the fourth octet would be available for subnetting and hosts. However, in supernetting, we can use bits in the usually reserved third octet to combine networks. With a 255.255.248.0 subnet, counting the zeros, we have three bits to work with. Each bit can be either a 1 or a 0. So, for each bit we get two choices. With three bits, it then becomes 2^3=8 subnets (don't subtract 2 in supernetting like you would in subnetting). Because the first network ID is 220.78.168.0, the router know to count up eight networks starting with 220.78.168.0 and going to 220.78.175.0.

Top Of Page

Summary

Well, we hope you've found all of this helpful and enlightening. To make sure none of it got lost in the shuffle, let's take a moment to recap things.

We explained the difference between static and dynamic IP routing, and also the host configuration requirements to communicate with a static or dynamic IP router. We configured a computer running Windows NT to function as an IP router. We also talked about building a static routing table and the use of the TRACERT utility, which allows us to isolate route or network link problems.

In regards to the NT TCP/IP 4.0 test, we hope you were paying attention in this chapter. Microsoft seems to want you to understand how routing works. So, to make sure you were paying attention, go through the exercises below.

Top Of Page

Review Questions

You have three Class C network addresses: 203.200.5.0, 203.200.6.0, and 203.200.7.0. You want to combine these addresses into one logical network to increase the host IDs that you can have on one subnet. Which subnet mask should you assign?

1. 255.255.252.0
2. 255.255.254.0
3. 255.255.255.252
4. 255.255.255.254

You can ping all the computers on your subnet and your default gateway, but you cannot ping any of the computers on a remote subnet. Other users on your subnet can ping computers on the remote subnet. What could be the problem?

1. The subnet mask on the router is invalid.
2. The subnet mask on your computer is invalid.
3. The default gateway address on your computer is invalid.
4. The computers on the remote subnet are not WINS enabled.
5. The route to the remote subnet has not been established on the router.

Using Windows NT Explorer, your Windows NT Workstation can connect to a remote server, but not to a server on your local subnet. What is most likely the cause of the problem?

1. Invalid default gateway address on your workstation
2. Invalid default gateway address on the local server
3. Invalid subnet mask on your workstation
4. Invalid subnet mask on the remote server

Using Windows NT Explorer, your Windows NT Workstation cannot connect to a local server, but all other users can. When you run Network Monitor, you notice that each time the workstation attempts to connect to the server, it broadcasts an ARP request for the default gateway. What is most likely the cause of the problem?

1. Invalid default gateway address on the workstation

2. Invalid subnet mask on the workstation

3. The workstation has a duplicate IP address.

4. The workstation is not configured to use WINS.

Your NT Workstation cannot connect to a remote server, but all other workstations can. The network is configured as follows: two subnets, one router. The router has two interfaces: Network A with 131.107.32.1 and Network B as 131.107.64.1. All computers use a subnet mask of 255.255.240.0. You are located on Network B. When you run IPCONFIG on your workstation, you receive the following output:

• IP Address	131.107.82.17
Subnet Mask	255.255.240.0
Default Gateway	131.108.64.1

What is most likely the cause of the problem?

1. IP address on workstation is invalid

2. Subnet mask on workstation is invalid

3. Default gateway on workstation is invalid

4. IP address on server is invalid

5. Default gateway on server is invalid

You are troubleshooting a Windows NT Server computer on a TCP/IP network. The server is located on a subnet with a network ID of 142.170.2.0. The default gateway address is 142.170.2.1. Users on a remote subnet cannot access the server. You run ipconfig /all at the server and receive the following output:

• Host Name	pdctest
DNS Servers	
Node Type	Hybrid
NetBIOS scope ID	
IP Routing Enabled	no
WINS Proxy Enabled	no
NetBIOS Resolution Uses DNS	no
Physical Address	00-20-AF-CA-E5-27

DHCP Enabled	no
IP Address	142.170.2.223
Subnet Mask	255.255.0.0
Default Gateway	142.170.2.1
Primary WINS Server	142.170.2.46

What is most likely the cause of the problem?

1. Incorrect subnet mask
2. NetBIOS scope ID is incorrect
3. NetBIOS node type is incorrect
4. IP address is out of range for this subnet

You have five Windows NT Server computers all configured as static routers. Which utility should you use to identify the path that a packet takes as it passes through the routers?

1. Network Monitor
2. ROUTE.EXE
3. TRACERT.EXE
4. IPCONFIG.EXE
5. NETSTAT.EXE

You have an NT Server on a remote subnet. You cannot ping this server by its IP address, but you can ping your default gateway and all other computers on the remote subnet. What are two likely causes of the problem?

1. The server is not WINS enabled.
2. The server has an invalid subnet mask.
3. The server has an invalid default gateway address.
4. Your workstation is configured with an invalid subnet mask.
5. Your workstation is configured with an invalid default gateway address.

You can connect through NT Explorer to all workstations on your local subnet, but cannot connect to any workstations or servers on a remote subnet. Which IP address should you ping first to diagnose the problem?

1. The local server
2. The default gateway
3. The remote server

4. None, TCP/IP isn't loaded

5. None, reboot the router

Your Windows NT Workstation cannot connect to a remote Windows NT Server using NT Explorer. All other computers can connect to the remote NT Server. You run Network Monitor and notice that each time the workstation attempts to connect to the server, the workstation broadcasts an ARP request for the remote server's IP address. What is most likely the cause of the problem?

1. The workstation is configured with an invalid default gateway address.

2. The workstation is configured with an invalid subnet mask.

3. The workstation is configured with a duplicate IP address.

4. The workstation is not configured to use WINS.

You have a RAS server that is connected to an Internet Service Provider via ISDN. You want your Windows 95 workstations to use the RAS server to access the Internet. How should the default gateway addresses be configured?

1. The default gateway address on the RAS server specifies the IP address of the ISP's router interface to your internal network.

2. The default gateway address on the RAS server specifies the IP address of the ISP's router interface to the Internet.

3. The default gateway address on each Windows 95 computer specifies the IP address of the ISP's router interface to your internal network.

4. The default gateway address on each Windows 95 machine specifies the IP address of the RAS server's network interface to your internal network.

Your NT Server has three network adapters. You configure the server to route TCP/IP packets and you want it to be able to automatically update its routing tables by using routing information from other routers on the network. Which service must you install on the Windows NT Server computer?

1. RIP for IP

2. RIP for NWLink IPX/SPX Compatible Transport

3. The DHCP Relay Agent

4. The DHCP Service

You have five multihomed Windows NT Server computers running TCP/IP and routing TCP/IP packets. You want to configure the routing tables on these servers with a minimum of administrative effort. What should you do?

1. Use NETSTAT.EXE to configure the routing tables.

2. Use ROUTE.EXE to configure the routing tables.

3. Install the DHCP Relay Agent.

4. Install RIP for IP.

You have a multihomed Windows NT Server that you want to configure as a TCP/IP static router. What two steps must you complete?

1. Enable IP forwarding.

2. Configure each network adapter with a unique subnet mask.
3. Configure each network adapter with an IP address, and ensure that each IP address is from a different subnet.
4. Configure each network adapter with an IP address, and ensure that each IP address is from a different address class.

You have installed two routers on one subnet to provide redundancy. When one router fails, users start complaining that they cannot access remote subnets, even though you have a second router on the network. What should you do to prevent this problem from occurring the next time a router fails?

1. Configure each workstation with multiple default gateway addresses.
2. Configure each workstation with multiple IP addresses.
3. Install a WINS server on each subnet.
4. Install the DHCP Relay Agent on both routers.

The network is configured as follows: two subnets, one router. The router has two interfaces: Network A with 131.107.32.1 and Network B with 131.107.64.1. All computers use a subnet mask of 255.255.224.0. Your NT Workstation cannot connect to a remote server on Network A, but all other workstations on Network B can. You are located on Network B. When you run IPCONFIG on your workstation, you receive the following output:

• IP Address	131.107.82.17
Subnet Mask	255.255.240.0
Default Gateway	31.108.32.1

What is most likely the cause of the problem?

1. IP address on workstation is invalid
2. Subnet mask on workstation is invalid
3. Default gateway on workstation is invalid
4. IP address on server is invalid
5. Default gateway on server is invalid

The network is configured as follows: two subnets, one router. The router has two interfaces: Network A with 131.107.32.1 and Network B with 131.107.64.1. All computers use a subnet mask of 255.255.224.0. Your NT Workstation cannot connect to a remote server on Network A, but all other workstations on Network B can. You are located on Network B. When you run IPCONFIG on your workstation, you receive the following output:

• IP Address	131.107.82.17
Subnet Mask	255.255.240.0
Default Gateway	131.108.64.1

What is most likely the cause of the problem?

1. IP address on workstation is invalid
2. Subnet mask on workstation is invalid
3. Default gateway on workstation is invalid
4. IP address on server is invalid
5. Default gateway on server is invalid

Situation: You are installing an NT TCP/IP server with three network adapters. You also plan to use this server as a router.

Required result:

- The new server must be configured to route TCP/IP.

Optional desired results:

- The server must dynamically update its routing tables.
- The server must provide IP addresses to all clients located on all subnets.
- The server must be able to send trap messages across the network to a Windows NT workstation computer.

Proposed solution:

- Install TCP/IP and configure one IP address for each of the servers network adapters.
- Install PPTP on the server.
- Install DHCP on the server and configure one scope for each subnet.
- Install SNMP on the server and configure SNMP to forward trap messages to the workstation.

Which results does the proposed solution produce?

1. The proposed solution produces the required result and produces all of the optional desired results.
2. The proposed solution produces the required result and produces only two of the optional desired results.
3. The proposed solution produces the required result but does not produce any of the optional desired results.
4. The proposed solution does not produce the required result.

Situation: You are installing an NT TCP/IP server with three network adapters. You also plan to use this server as a router.

Required result:

- The new server must be configured to route TCP/IP.

Optional desired results:

- The server must dynamically update its routing tables.
- The server must provide IP addresses to all clients located on all subnets.
- The server must be able to send trap messages across the network to a Windows NT workstation computer.

Proposed solution:

- Install TCP/IP and configure one IP address for each server's network adapters.
- Enable IP forwarding on the server.
- Install PPTP on the server.
- Install the DHCP Relay Agent on the server.

Which results does the proposed solution produce?

1. The proposed solution produces the required result and produces all of the optional desired results.
2. The proposed solution produces the required result and produces only two of the optional desired results.
3. The proposed solution produces the required result but does not produce any of the optional desired results.
4. The proposed solution does not produce the required result.

Situation: You are installing an NT TCP/IP server with three network adapters. You also plan to use this server as a router.

Required result:

- The new server must be configured to route TCP/IP.

Optional desired results:

- The server must dynamically update its routing tables.
- The server must provide IP addresses to all clients located on all subnets.
- The server must be able to send trap messages across the network to a Windows NT workstation computer.

Proposed solution:

- Install TCP/IP and configure one IP address for each of the servers network adapters.
- Enable IP forwarding on the server.
- Install RIP for IP on the server.
- Install DHCP with scopes for all subnets.

Which results does the proposed solution produce?

1. The proposed solution produces the required result and produces all of the optional desired results.

2. The proposed solution produces the required result and produces only two of the optional desired results.
3. The proposed solution produces the required result but does not produce any of the optional desired results.
4. The proposed solution does not produce the required result.

Situation: You are installing an NT TCP/IP server with three network adapters. You also plan to use this server as a router.

Required result:

- The new server must be configured to route TCP/IP.

Optional desired results:

- The server must dynamically update its routing tables.
- The server must provide IP addresses to all clients located on all subnets.
- The server must be able to send trap messages across the network to a Windows NT workstation computer.

Proposed solution:

- Install TCP/IP and configure one IP address for the server's network adapters.
- Enable IP forwarding on the server.
- Install DHCP with scopes for all subnets.
- Install SNMP on the server and configure SNMP to forward trap messages to the workstation.
- Install a third-party SNMP manager on the server.

Which results does the proposed solution produce?

1. The proposed solution produces the required result and produces all of the optional desired results.
2. The proposed solution produces the required result and produces only two of the optional desired results.
3. The proposed solution produces the required result but does not produce any of the optional desired results.
4. The proposed solution does not produce the required result.

Your boss frantically comes up to you and says he put two NIC cards in his NT Server, but the workstations on each segment can't see each other. He knows that to route IP packets to other networks each multihomed computer (static router) must be configured two ways, but he can't remember what they are. What are the two things you need to set on the NT Server?

You get a call from a company that thinks they need some routers. They have a small network and read in a magazine that they should employ static routing. They want to know more about how static routing would meet their networking needs. What do you tell them?

It's your first day on the job at Terrific Technology Teaching Center, and as a co-instructor you are asked to teach the enabling of IP routers. Take a moment to explain the procedure now.

Later, a confused student comes up to you and asks if she needs to add a routing table to a computer running as a multihomed computer and connecting two subnet segments. What do you tell her, and why?

You have two NT servers and a router to the Internet. Should you build a static router between the NT servers, or will the dynamic router to the Internet be sufficient?

What information would you put into the static routing table?

Using supernetting, assign the missing IP and subnet mask values for each customer below:

• Beginning IP address	192.24.0.1
Ending IP address	192.24.7.8
Subnet Mask	
Beginning IP address	
Ending IP address	192.34.31.254
Subnet Mask	255.255.240.0
Beginning IP address	192.24.8.1
Ending IP address	
Subnet Mask	255.255.252.0
Beginning IP address	192.24.14.1
Ending IP address	192.24.15.254
Subnet Mask	

About the Authors

Todd Lammle is a Microsoft Certified Trainer (MCT) with over fifteen years of experience with LANs and WANs. He is president of GlobalNet Systems, a network integration firm in Colorado.

Monica Lammle is a Microsoft Certified Product Specialist (MCPS) in TCP/IP.

James Chellis, a Microsoft Certified Professional (MCP), is president of EdgeTek Computer Education, a national network training company and Microsoft Solution Provider.

We at Microsoft Corporation hope that the information in this work is valuable to you. Your use of the information contained in this work, however, is at your sole risk. All information in this work is provided "as - is", without any warranty, whether express or implied, of its accuracy, completeness, fitness for a particular purpose, title or non-infringement, and none of the third-party products or information mentioned in the work are authored, recommended, supported or guaranteed by Microsoft Corporation. Microsoft Corporation shall not be liable for any damages you may sustain by using this information, whether direct, indirect, special, incidental or consequential, even if it has been advised of the possibility of such damages. All prices for products mentioned in this document are subject to change without notice. International rights = English only.

International rights = English only.

Top Of Page

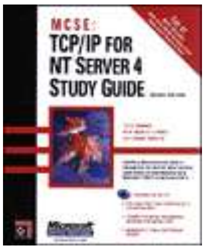

Top Of Page

Configuring and Troubleshooting DDR Backup

Contents

Before Calling the Cisco Systems TAC Team
Related Information

Introduction

Dialup is simply the application of the public switched telephone network (PSTN) that carries data on behalf of the end user. It involves a customer premises equipment (CPE) device sending the telephone switch a phone number to which to direct a connection. The Cisco3600, AS5200, AS5300, and AS5800 are all examples of routers that have the ability to run a PRI along with banks of digital modems. The AS2511, on the other hand, is an example of a router that communicates with external modems.

Prerequisites

Requirements

Readers of this document should be knowledgeable of the following:

The carrier market has grown significantly, and the market now demands higher modem densities. The answer to this need is a higher degree of interoperation with the telephone company equipment and the development of the digital modem. This is a modem that is capable of direct digital access to the PSTN. As a result, faster CPE modems have now been developed that take advantage of the clarity of signal that the digital modems enjoy. The fact that the digital modems connecting into the PSTN through a PRI or BRI can transmit data at over 53k using the V.90 communication standard, attests to the success of the idea.

The first access servers were the Cisco2509 and Cisco2511. The AS2509 could support 8 incoming connections using external modems, and the AS2511 could support 16. The AS5200 was introduced with 2 PRIs and could support 48 users using digital modems, and it represented a major leap forward in technology. Modem densities have increased steadily with the AS5300 supporting 4 and then 8 PRIs. Finally, the AS5800 was introduced to fill the needs of carrier class installations needing to handle dozens of incoming T1s and hundreds of user connections.

A couple of outdated technologies bear mentioning in a historical discussion of dialer technology. 56Kflex is an older (pre-V.90) 56k modem standard that was proposed by Rockwell. Cisco supports version 1.1 of the 56Kflex standard on its internal modems, but recommends migrating the CPE modems to V.90 as soon as possible. Another outdated technology is the AS5100. The AS5100 was a joint venture between Cisco and a modem manufacturer. The AS5100 was created as a way to increase modem density through the use of quad modem cards. It involved a group of AS2511s built as cards that inserted into a backplane shared by quad modem cards, and a dual T1 card.

Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

Troubleshooting Incoming Calls

Troubleshooting an incoming call starts at the bottom and works its way up. The general flow of reasoning looks for the following:

Do we see the call arrive? (A yes answer advances to the next question)

Does the receiving end answer the call?

Does the call complete?

Is data passing across the link?

Is the session established? (PPP or terminal)

For modem connections, a data call looks the same as a terminal session coming in until the end where the data call goes to negotiate PPP.

For incoming calls involving digital modems, first make sure the underlying ISDN or CAS is receiving the call. If using an external modem, the ISDN and CAS group sections can be skipped.

Incoming ISDN Call Troubleshooting

Use the command debug isdn q931. Here's an example output from a successful connection:

```
Router# debug isdn q931

RX <- SETUP pd = 8 callref = 0x06

 Bearer Capability i = 0x8890

 Channel ID i = 0x89

 Calling Party Number i = 0x0083, `5551234'

TX -> CONNECT pd = 8 callref = 0x86

RX <- CONNECT_ACK pd = 8 callref = 0x06
```

The setup message indicates that a connection is being initiated by the remote end. The call reference numbers are maintained as a pair. In this case the call reference number for the incoming side of the connection is 0x06, and the call reference number of the outbound side of the connection is 0x86. The Bearer Capability (often referred to as the bearercap) tells the router what kind of call is coming in. In this case the connection is type 0x8890. That value indicates "ISDN Speed 64 Kb/s". If the bearercap had been 0x8090A2, it would have indicated "Speech/voice call u-law".

If no setup message came in, you should verify the correct number by calling it manually, if it is voice provisioned. You should also check the status of the ISDN interface (refer to Using the show isdn status Command for BRI Troubleshooting). If that all checks out, make sure that the call originator is making the correct call. This can be done by contacting the telephone company. The call originator can trace the call to see where it?s being sent. If the connection is long distance, try a different long distance carrier using a 1010 long distance code.

If the call coming in is an async modem call, make sure the line is provisioned to allow voice calls.

Note: BRI async modem calling is a feature of 3600 routers running 12.0(3)T, or later. It requires a recent hardware revision of the BRI interface network module. WIC modules do not support async modem calling.

If the call arrived but did not complete, look for a cause code (see Table 17-10). A successful completion is indicated by connect-ack.

If this is an async modem call, move forward to the "Incoming Modem Call Troubleshooting" section.

At this point the ISDN call is connected, but no data has been seen coming across the link. Use the command debug ppp negotiate to see if any PPP traffic is coming across the line. If you do not see traffic, there may be a speed mismatch. To determine if this is the case, use the show running-config privileged exec command to view the router configuration. Check the dialer map interface configuration command entries in the local and remote router. These entries should look similar to the following:

```
dialer map ip 131.108.2.5 speed 56 name C4000
```

For dialer profiles, a map-class needs to be defined in order to set the speed. Note that, by default, ISDN interfaces attempt to use 64K communications speeds on each channel.

For detailed information on configuring dialer maps and profiles, refer to the Cisco IOS Dial Solutions Configuration Guide, Dial Solutions Command Reference, and the Dial Solutions Quick Configuration Guide.

If you receive valid PPP packets, the link is up and working. You should proceed to the "Troubleshooting PPP" section at this time.

Incoming CAS Call Troubleshooting

To troubleshoot the CAS group serving connectivity to the modems, use the commands debug modem, debug modem csm, and debug cas.

Note: The debug cas command first appeared in 12.0(7)T for the AS5200 and AS5300. Earlier versions of IOS use the system level configuration command service internal along with the exec command modem-mgmt debug rbs. Debugging this information on an AS5800 requires connecting to the trunk card itself.

First, determine if the telephone company switch went "offhook" to signal the incoming call. If it did not, verify the number being called. Do this by attaching a phone to the originating side's phone line and calling the number. If the call comes in properly, the problem is in the originating CPE. If the call still does not show up on the CAS, check the T1 (chapter 15).In this instance, use the debug serial interfaces command.

The following shows a good connection using debug modem CSM:

```
Router# debug modem csm

CSM_MODEM_ALLOCATE: slot 1 and port 0 is allocated.

MODEM_REPORT(0001): DEV_INCALL at slot 1 and port 0

CSM_PROC_IDLE: CSM_EVENT_ISDN_CALL at slot 1, port 0

CSM_RING_INDICATION_PROC: RI is on

CSM_RING_INDICATION_PROC: RI is off

CSM_PROC_IC1_RING: CSM_EVENT_MODEM_OFFHOOK at slot 1, port 0

MODEM_REPORT(0001): DEV_CONNECTED at slot 1 and port 0

CSM_PROC_IC2_WAIT_FOR_CARRIER: CSM_EVENT_ISDN_CONNECTED at slot 1, port 0
```

In this example, the call was directed to a modem. If your call was directed to a modem, proceed to the "Incoming Modem Call Troubleshooting" section, below.

Incoming Modem Call Troubleshooting

Use the following debug commands when troubleshooting incoming modem calls:

```
debug modem

debug modem csm (for integrated digital modems)
```

Use the following debug commands in conjunction to indicate the new call coming in:

```
debug isdn q931

debug cas
```

Assuming the call reaches the modem, the modem needs to pick the call up.

Tips for debugging External Modems

To facilitate debugging on an external modem connected to a TTY line, increase the speaker volume. This helps to make some problems more apparent.

When the originating modem calls, does the receiving modem ring? If not, verify the number and try a manual call from the remote site. Try using a regular phone on the receiving end as well. Replace cables and hardware as needed.

Async Modem Call Pickup

If an external modem is not answering, check the cabling between the modem and the access server or router. Confirm that the modem is connected to the TTY or auxiliary port on the router with a rolled RJ-45 cable and an MMOD DB-25 adapter. Cisco recommends and supports this cable configuration for RJ-45 ports. Note that these connectors are typically labeled: Modem.

RJ-45 cabling comes in a few types: straight, rolled, and crossover. You can determine the cabling type by holding the two ends of an RJ-45 cable side-by-side. You'll see eight colored strips, or pins, at each end.

If the order of the colored pins is the same at each end, the cable is straight.

If the order of the colors is reversed at each end, the cable is rolled.

The cable is a crossover cable if colors indicate the following:

RJ45 to RJ45 crossover cable:

RJ45 RJ45

5 ------------------ 2

2 ------------------ 5

4 ------------------ 1

1 ------------------ 4

To make sure the signaling is OK, use the show line command outlined in chapter 16.

Cabling issues aside, an external modem needs to be initialized to auto-answer. Check the remote modem to see whether it is set to auto-answer. Usually, an AA indicator light is on when auto-answer is set. Set the remote modem to auto-answer if it is not already set. For information on verifying and changing the modem's settings, refer to your modem documentation. Use a reverse telnet to initialize the modem (refer to chapter 16).

Digital (Integrated) Modem Call Pickup

On an external modem it is clear whether the call is being answered, but internal modems require a manual call to the receiving number. Listen for the answer back tone (ABT). If you do not hear an ABT, check the configuration for the following two things:

Make sure the command isdn incoming-voice modem exists under any ISDN interfaces handling incoming modem connections.

Under the line configuration for the modem's TTY, make sure the command modem inout exists.

It is also possible that the Call Switching Module (CSM) did not allocate an internal modem to handle the incoming call. This problem can be caused by modem or resource pools being configured for too few incoming connections. It may also mean that the access server may simply be out of modems. Check the availability of modems and adjust the modem pool or resource pool manager settings appropriately. If a modem was allocated and the configuration shows modem inout, gather debugs and contact Cisco for assistance.

Modem Trainup

If the receiving modem raises DSR, the trainup was successful. Trainup failures can indicate a circuit problem or modem incompatibility.

To get to the bottom of an individual modem problem, go to the AT prompt at the originating modem while it's attached to the POTS line of interest. If calling into a digital modem in a Cisco access server, be prepared to record a .wav file of the trainup music, or digital impairment learning sequence (DIL). The DIL is the musical score (PCM sequence) that the originating V.90 analog modem tells the receiving digital modem to play back. The sequence allows the analog modem to discern any digital impairment in the circuit; such as multiple D/A conversions, a law/u-law, robbed bits, or digital pads. If you don't hear the DIL, the modems did not negotiate V.90 in V.8/V.8bis (that is., a modem compatibility issue). If you do hear the DIL and a retrain in V.34, the analog modem decided (on the basis of the DIL playback) that V.90 was infeasible.

Does the music have noise in it? If so, then clean up the circuit.

Does the client give up quickly, without running V.34 training? For example, perhaps it doesn't know what to do when it hears V.8bis. In this case you should try disabling V.8bis (hence K56Flex) on the server (if acceptable). You should get new client firmware or swap out the client modem. Alternately, the dialing end could insert five commas at the end of the dial string. This delays the calling modem's listen and will cause the V.8bis tone from the receiving server to timeout without affecting the client modem. Five commas in the dial string is a general guideline and might need adjusting to allow for local conditions.

Session Establishment

At this point in the sequence, the modems are connected and trained up. Now it's time to find out if any traffic is coming across properly.

If the line receiving the call is configured with autoselect ppp and the async interface is configured with async mode interactive, use the command debug modem to verify the autoselect process. As traffic comes in over the async link, the access server will examine the traffic to determine whether the traffic is character-based or packet-based. Depending on the determination, the access server will then either start a PPP session or go no farther than having an exec session on the line.

A normal autoselect sequence with inbound PPP LCP packets:

```
*Mar  1 21:34:56.958: TTY1: DSR came up
```

```
*Mar  1 21:34:56.962: tty1: Modem: IDLE->READY
*Mar  1 21:34:56.970: TTY1: EXEC creation
*Mar  1 21:34:56.978: TTY1: set timer type 10, 30 seconds
*Mar  1 21:34:59.722: TTY1: Autoselect(2) sample 7E

!--- The inbound traffic is displayed in hexadecimal format. This is based on the
!--- bits coming in over the line, regardless of whether the bits are ASCII
!--- characters or elements of a packet. The bits represented in this example are
!--- correct for a LCP packet. Anything different would be either a malformed packet
!--- or character traffic.

*Mar  1 21:34:59.726: TTY1: Autoselect(2) sample 7EFF
*Mar  1 21:34:59.730: TTY1: Autoselect(2) sample 7EFF7D
*Mar  1 21:34:59.730: TTY1: Autoselect(2) sample 7EFF7D23
*Mar  1 21:34:59.734: TTY1 Autoselect cmd: ppp negotiate

!--- Having determined that the inbound traffic is actually an LCP packet, the access
!--- server triggers the PPP negotiation process.

*Mar  1 21:34:59.746: TTY1: EXEC creation
*Mar  1 21:34:59.746: TTY1: create timer type 1, 600 seconds
*Mar  1 21:34:59.794: TTY1: destroy timer type 1 (OK)
*Mar  1 21:34:59.794: TTY1: destroy timer type 0
*Mar  1 21:35:01.798: %LINK-3-UPDOWN: Interface Async1, changed state to up

!--- The async interface changes state to up, and the PPP negotiation (not shown)
!--- commences.
```

If the call is a PPP session and if async mode dedicated is configured on the async interface, use the command debug ppp negotiation to see if any configuration request packets are coming from the remote end. The debugs show these as CONFREQ. If you observe both inbound and outbound PPP packets, proceed to "Troubleshooting PPP". Otherwise, connect from the call-originating end with a character-mode (or "exec") session (that is, a non-PPP session).

Note: If the receiving end displays async modem dedicated under the async interface, an exec dial-in only shows what appears to be random ASCII garbage. To allow a terminal session and still have PPP capability, use the async interface configuration command async mode interactive. Under the associated line's configuration, use the command autoselect ppp.

Modem Cannot Send or Receive Data

If the modems connect with a terminal session and no data comes across, check the following possible causes and suggested courses of action:

Modem speed setting is not locked

Use the show line exec command on the access server or router. The output for the auxiliary port should indicate the currently configured Tx and Rx speeds.

For an explanation of the output of the show line command, see the "Using Debug Commands" section in chapter 15.

If the line is not configured to the correct speed, use the speed line configuration command to set the line speed on the access server or router line. Set the value to the highest speed in common between the modem and the access server or router port. To set the terminal baud rate, use the speed line configuration command. This command sets both the transmit (to terminal) and receive (from terminal) speeds.

Syntax:

```
speed bps
```

Syntax Description:

bps - Baud rate in bits per second (bps). The default is 9600 bps.

The following example sets lines 1 and 2 on a Cisco 2509 access server to 115200 bps:

```
line 1 2

speed 115200
```

Note: If, for some reason, you cannot use flow control, limit the line speed to 9600 bps. Faster speeds are likely to result in lost data.

Use the show line exec command again and confirm that the line speed is set to the desired value.

When you are certain that the access server or router line is configured for the desired speed, initiate a reverse Telnet session to the modem via that line. For more information, see the section "Establishing a Reverse Telnet Session to a Modem" in chapter 16.

Use a modem command string that includes the "lock DTE speed" command for your modem. See your modem documentation for exact configuration command syntax.

Note: The lock DTE speed command, which might also be referred to as port rate adjust or buffered mode, is often related to the way in which the modem handles error correction. This command varies widely from one modem to another.

Locking the modem speed ensures that the modem always communicates with the Cisco access server or router at the speed configured on the Cisco auxiliary port. If this command is not used, the modem reverts to the speed of the data link (the telephone line), instead of communicating at the speed configured on the access server.

Hardware flow control not configured on local or remote modem or router

Use the show line aux-line-number exec command and look for the following in the Capabilities field:

Capabilities: Hardware Flowcontrol In, Hardware Flowcontrol Out

For more information, refer to Interpreting Show Line Output in Chapter 16.

If there is no mention of hardware flow control in this field, hardware flow control is not enabled on the line. Hardware flow control for access server-to-modem connections is recommended.

For an explanation of the output of the show line command, see the section "Using Debug Commands" in chapter 15.

Configure hardware flow control on the line using the flowcontrol hardware line configuration command.

To set the method of data flow control between the terminal or other serial device and the router, use the flowcontrol line configuration command. Use the no form of this command to disable flow control.

Syntax:

flowcontrol {none | software [lock] [in | out] | hardware [in | out]}

Syntax Description:

none - Turns off flow control.

software - Sets software flow control. An optional keyword specifies the direction: in causes the Cisco IOS software to listen to flow control from the attached device, and out causes the software to send flow control information to the attached device. If you do not specify a direction, both are assumed.

lock - Makes it impossible to turn off flow control from the remote host when the connected device needs software flow control. This option applies to connections using the Telnet or rlogin protocols.

hardware - Sets hardware flow control. An optional keyword specifies the direction: in causes the software to listen to flow control from the attached device, and out causes the software to send flow control information to the attached device. If you do not specify a direction, both are assumed. For more information about hardware flow control, see the hardware manual that was shipped with your router.

Example:

The following example sets hardware flow control on line 7:

```
line 7

flowcontrol hardware
```

Note: If for some reason you cannot use flow control, limit the line speed to 9600 bps. Faster speeds are likely to result in lost data.

After enabling hardware flow control on the access server or router line, initiate a reverse Telnet session to the modem via that line. For more information, see the section "Establishing a Reverse Telnet Session to a Modem" in chapter 16.

Use a modem command string that includes the RTS/CTS Flow command for your modem. This command ensures that the modem is using the same method of flow control (that is, hardware flow control) as the Cisco access server or router. See your modem documentation for exact configuration command syntax.

Misconfigured dialer map commands

Use the show running-config privileged exec command to view the router configuration. Check the dialer map command entries to see whether the broadcast keyword is specified.

If the keyword is missing, add it to the configuration.

Syntax:

```
dialer map protocol next-hop-address [name hostname] [broadcast] [dial-string]
```

Syntax Description:

protocol - The protocol subject to mapping. Options include IP, IPX, bridge, and snapshot.

next-hop-address - The protocol address of the opposite site's async interface.

name hostname - A required parameter used in PPP authentication. It is the name of the remote site for which the dialer map is created. The name is case sensitive and must match the hostname of the remote router.

broadcast - An optional keyword that broadcast packets (for example, IP RIP or IPX RIP/SAP updates) that is forwarded to the remote destination. In static routing sample configurations, routing updates are not desired and the broadcast keyword is omitted.

dial-string - The remote site's phone number. Any access codes (for example, 9 to get out of an office, international dialing codes, area codes) must be included.

Make sure that dialer map commands specify the correct next hop addresses.

If the next hop address is incorrect, change it using the dialer map command.

Make sure all other options in dialer map commands are correctly specified for the protocol you are using.

For detailed information on configuring dialer maps, refer to the Cisco IOS Wide-Area Networking Configuration Guide and Wide-Area Networking Command Reference.

Problem with dialing modem

Make sure that the dialing modem is operational and is securely connected to the correct port. Determine if another modem works when connected to the same port.

Debugging an incoming exec session generally falls into a few main categories:

Dialup client receives No exec Prompt

Dialup Session Sees "Garbage"

Dialup Session Opens in Existing Session

Dialup Receiving Modem Does Not Disconnect Properly

Dialup Client Receives No exec Prompt

Autoselect is enabled on the line

Attempt to access exec mode by pressing Enter.

Line is configured with the no exec command

Use the show line exec command to view the status of the appropriate line.

Check the Capabilities field to see if it says "exec suppressed." If this is the case, the no exec line configuration command is enabled.

Configure the exec line configuration command on the line to allow exec sessions to be initiated. This command has no arguments or keywords.

The following example turns on the exec on line 7:

```
line 7
```

exec

Flow control is not enabled.

or

Flow control is enabled only on one device (either DTE or DCE).

or

Flow control is misconfigured.

Use the show line aux-line-number exec command and look for the following in the Capabilities field:

```
Capabilities: Hardware Flowcontrol In, Hardware Flowcontrol Out
```

For more information, refer to Interpreting Show Line Output in Chapter 16.

If there is no mention of hardware flow control in this field, hardware flow control is not enabled on the line. Hardware flow control for access server-to-modem connections is recommended.

For an explanation of the output from the show line command, see the "Using Debug Commands" section in chapter 15.

Configure hardware flow control on the line using the flowcontrol hardware line configuration command. The following example sets hardware flow control on line 7:

```
line 7

flowcontrol hardware
```

Note: If for some reason you cannot use flow control, limit the line speed to 9600 bps. Faster speeds are likely to result in lost data.

After enabling hardware flow control on the access server or router line, initiate a reverse Telnet session to the modem via that line. For more information, see the section "Establishing a Reverse Telnet Session to a Modem" in chapter 16.

Use a modem command string that includes the RTS/CTS Flow command for your modem. This command ensures that the modem is using the same method of flow control (that is, hardware flow control) as the Cisco access server or router. See your modem documentation for exact configuration command syntax.

Modem speed setting is not locked

Use the show line exec command on the access server or router. The output for the auxiliary port should indicate the currently configured Tx and Rx speeds.

For an explanation of the output of the show line command, see the "Using Debug Commands" section in chapter 15.

If the line is not configured to the correct speed, use the speed line configuration command to set the line speed on the access server or router line. Set the value to the highest speed in common between the modem and the access server or router port. To set the terminal baud rate, use the speed line configuration command. This command sets both the transmit (to terminal) and receive (from terminal) speeds.

Syntax:

speed bps

Syntax Description:

bps - Baud rate in bits per second (bps). The default is 9600 bps.

Example:

The following example sets lines 1 and 2 on a Cisco 2509 access server to 115200 bps:

line 1 2

speed 115200

Note: If for some reason you cannot use flow control, limit the line speed to 9600 bps. Faster speeds are likely to result in lost data.

Use the show line exec command again and confirm that the line speed is set to the desired value.

When you are certain that the access server or router line is configured for the desired speed, initiate a reverse Telnet session to the modem via that line. For more information, see the section "Establishing a Reverse Telnet Session to a Modem" in chapter 16.

Use a modem command string that includes the lock DTE speed command for your modem. See your modem documentation for exact configuration command syntax.

Note: The lock DTE speed command, which might also be referred to as port rate adjust or buffered mode, is often related to the way in which the modem handles error correction. This command varies widely from one modem to another.

Locking the modem speed ensures that the modem always communicates with the Cisco access server or router at the speed configured on the Cisco auxiliary port. If this command is not used, the modem reverts to the speed of the data link (the telephone line) instead of communicating at the speed configured on the access server.

Dialup Sessions Sees "Garbage"

Modem speed setting is not locked

Use the show line exec command on the access server or router. The output for the auxiliary port should indicate the currently configured Tx and Rx speeds.

For an explanation of the output of the show line command, see the "Using Debug Commands" section in chapter 15.

If the line is not configured to the correct speed, use the speed line configuration command to set the line speed on the access server or router line. Set the value to the highest speed in common between the modem and the access server or router port.

To set the terminal baud rate, use the speed line configuration command. This command sets both the transmit (to terminal) and receive (from terminal) speeds.

Syntax:

```
speed bps
```

Syntax Description:

bps Baud rate in bits per second (bps). The default is 9600 bps.

Example:

The following example sets lines 1 and 2 on a Cisco 2509 access server to 115200 bps:

```
line 1 2

speed 115200
```

Note: If for some reason you cannot use flow control, limit the line speed to 9600 bps. Faster speeds are likely to result in lost data.

Use the show line exec command again and confirm that the line speed is set to the desired value.

When you are certain that the access server or router line is configured for the desired speed, initiate a reverse Telnet session to the modem via that line. For more information, see the section "Establishing a Reverse Telnet Session to a Modem" in chapter 16.

Use a modem command string that includes the lock DTE speed command for your modem. See your modem documentation for exact configuration command syntax.

Note: The lock DTE speed command, which might also be referred to as port rate adjust or buffered mode, is often related to the way in which the modem handles error correction. This command varies widely from one modem to another.

Locking the modem speed ensures that the modem always communicates with the Cisco access server or router at the speed configured on the Cisco auxiliary port. If this command is not used, the modem reverts to the speed of the data link (the telephone line) instead of communicating at the speed configured on the access server.

Symptom: Remote dialin session opens up in an already-existing session initiated by another user. That is, instead of getting a login prompt, a dialin user sees a session established by another user (which might be a UNIX command prompt, a text editor session, and so forth).

Dialup Session Opens in Existing Session

Modem configured for DCD always high

The modem should be reconfigured to have DCD high only on CD. This is usually accomplished by using the &C1 modem command string, but check your modem documentation for the exact syntax for your modem.

You might have to configure the access server line to which the modem is connected with the no exec line configuration command. Clear the line with the clear line privileged exec command, initiate a reverse Telnet session with the modem, and reconfigure the modem so that DCD is high only on CD.

End the Telnet session by entering disconnect and reconfigure the access server line with the exec line configuration command

Modem control is not enabled on the access server or router

Use the show line exec command on the access server or router. The output for the auxiliary port should be show inout or RIisCD in the Modem column. This indicates that modem control is enabled on the line of the access server or router.

For an explanation of the show line output, see the "Using Debug Commands" section in chapter 15.

Configure the line for modem control using the modem inout line configuration command. Modem control is now enabled on the access server.

Note: Be certain to use the modem inout command instead of the modem dialin command while the connectivity of the modem is in question. The latter command allows the line to accept incoming calls only. Outgoing calls will be refused, making it impossible to establish a Telnet session with the modem to configure it. If you want to enable the modem dialin command, do so only after you are certain the modem is functioning correctly.

Incorrect cabling

Check the cabling between the modem and the access server or router. Confirm that the modem is connected to the auxiliary port on the access server or router with a rolled RJ-45 cable and an MMOD DB-25 adapter. This cabling configuration is recommended and supported by Cisco for RJ-45 ports. These connectors are typically labeled: Modem.

There are two types of RJ-45 cabling: straight and rolled. If you hold the two ends of an RJ-45 cable side-by-side, you'll see eight colored strips, or pins, at each end. If the order of the colored pins is the same at each end, then the cable is straight. If the order of the colors is reversed at each end, then the cable is rolled.

The rolled cable (CAB-500RJ) is standard with Cisco's 2500/CS500.

Use the show line exec command to verify that the cabling is correct. See the explanation of the show line command output in the section "Using Debug Commands" in this chapter 15.

Dialup Receiving Modem Does Not Disconnect Properly

Modem is not sensing DTR

Enter the Hangup DTR modem command string. This command tells the modem to drop carrier when the DTR signal is no longer being received.

On a Hayes-compatible modem the &D3 string is commonly used to configure Hangup DTR on the modem. For the exact syntax of this command, see the documentation for your modem.

Modem control is not enabled on the router or access server

Use the show line exec command on the access server or router. The output for the auxiliary port should show inout or RIisCD in the Modem column. This indicates that modem control is enabled on the line of the access server or router.

For an explanation of the show line output, see the "Using Debug Commands" section in chapter 15.

Configure the line for modem control using the modem inout line configuration command. Modem control is now enabled on the access server.

Note: Be certain to use the modem inout command instead of the modem dialin command while the connectivity of the modem is in question. The latter command allows the line to accept incoming calls only. Outgoing calls will be refused, making it impossible to establish a Telnet session with the modem to configure it. If you want to enable the modem dialin command, do so only after you are certain the modem is functioning correctly.

Troubleshooting Outbound Calls

While the troubleshooting approach for incoming calls starts at the bottom, troubleshooting an outbound connection starts at the top. The general flow of reasoning looks for the following:

Does the Dial on Demand Routing (DDR) initiate a call? (A yes answer advances to the next question)

If this is an async modem, do the chat scripts issue the expected commands?

Does the call make it out to the PSTN?

Does the remote end answer the call?

Does the call complete?

Is data passing over the link?

Is the session established? (PPP or Terminal)

Verifying Dialer Operation

To see if the dialer is trying to make a call to its remote destination, use the command debug dialer events. More detailed information can be gained from debug dialer packet, but the debug dialer packet command is resource intensive and should not be used on a busy system that has multiple dialer interfaces operating.

The following line of debug dialer events output for an IP packet lists the name of the DDR interface and the source and destination addresses of the packet:

Dialing cause: Async1: ip (s=172.16.1.111 d=172.16.2.22)

If traffic does not initiate a dial attempt, the most common reason is improper configuration (either of the interesting traffic definitions, the state of the dialer interface, or the routing).

Traffic Does Not Initiate a Dial Attempt

Missing or incorrect "interesting traffic" definitions

Using the command show running-config, ensure that the interface is configured with a dialer-group and that there is a global level dialer-list configured with a matching number.

Ensure that the dialer-list command is configured to permit either an entire protocol or to permit traffic matching an access list

Verify that the access-list declares packets going across the link to be interesting. One useful test is to use the privileged exec command debug ip packet [list number] using the number of the pertinent access list. Then attempt to ping, or otherwise send traffic, across the link. If the interesting traffic filters have been properly defined, you will see the packets in the debug output. If there is no debug output from this test, then the access-list is not matching the packets.

Interface state

Use the command show interfaces [interface name] to ensure that the interface is in the state "up/up (spoofing)".

Interface in "standby" mode

Another (primary) interface on the router has been configured to use the dialer interface as a backup interface. Furthermore, the primary interface is not in a state of "down/down", which is required to bring the dialer interface out of standby mode. Also, a backup delay must be configured on the primary interface, or the backup interface command will never be enforced.

To check that the dialer interface will change from "standby" to "up/up (spoofing)", it is usually necessary to pull the cable from the primary interface. Simply shutting down the primary interface with the configuration command shutdown will not put the primary interface into "down/down", but instead will put it into "administratively down"-not the same thing.

In addition, if the primary connection is via Frame Relay, the Frame Relay configuration must be done on a point-to-point Serial sub-interface, and the telephone company must be passing the "Active" bit. This practice is also known as "end-to-end LMI".

Interface is "administratively down"

The dialer interface has been configured with the command shutdown. This is also the default state of any interface when a Cisco router is booted for the very first time. Use the interface configuration command no shutdown to remove this impediment.

Incorrect routing

Issue the exec command show ip route [a.b.c.d], where a.b.c.d is the address of the dialer interface of the remote router. If ip unnumbered is used on the remote router, use the address of the interface listed in the ip unnumbered command.

The output should show a route to the remote address via the dialer interface. If there is no route, ensure that static or floating static routes have been configured by examining the output of show running-config.

If there is a route via an interface other than the dialer interface, the implication is that DDR is being used as a backup. Examine the router configuration to make sure that static or floating static routes have been configured. The surest way to test the routing, in this case, is to disable the primary connection and execute the show ip route [a.b.c.d] command to verify that the proper route has been installed in the routing table.

Note: If you attempt this during live network operations, a dial event may be triggered. This sort of testing is best accomplished during scheduled maintenance cycles.

Placing the Call

If the routing and the interesting traffic filters are correct, a call should be initiated. This can be seen by using debug dialer events:

```
Async1 DDR: Dialing cause ip (s=10.0.0.1, d=10.0.0.2)

Async1 DDR: Attempting to dial 5551212
```

If the dialing cause is seen but no attempt is made to dial, the usual reason is a misconfigured dialer map or dialer profile.

Call Not Placed

Some possible problems and suggested actions are listed below:

Misconfigured dialer map

Use the command show running-config to ensure that the dialing interface is configured with at least one dialer map statement which points to the protocol address and called number of the remote site.

Misconfigured dialer profile

Use the command show running-config to ensure that the Dialer interface is configured with a dialer pool X command and that a dialer interface on the router is configured with a matching dialer pool-member X. If dialer profiles are not properly configured, you may see a debug message like:

```
Dialer1: Can't place call, no dialer pool set
```

Make sure that a dialer string is configured.

Async Outbound Calling - Verify Chat Script Operation

If the outbound call is a modem call, a chat script must execute in order for the call to proceed. For dialer map-based DDR, the chat script is invoked by the modem-script parameter in a dialer map command. If the DDR is dialer profile-based, this is accomplished by the command script dialer, configured on the TTY line. Both uses rely on a chat script existing in the router's global configuration, for example:

```
chat-script callout AT OK atdt\T TIMEOUT 60 CONNECT \c
```

In either event, the command to view the chat script activity is debug chat. If the dial string (that is, phone number) used in the dialer map or dialer string command were 5551212, the debug output would look like the following:

```
CHAT1: Attempting async line dialer script


CHAT1: Dialing using Modem script: callout & System script: none

CHAT1: process started

CHAT1: Asserting DTR

CHAT1: Chat script callout started

CHAT1: Sending string: AT

CHAT1: Expecting string: OK

CHAT1: Completed match for expect: OK

CHAT1: Sending string: atdt5551212

CHAT1: Expecting string: CONNECT

CHAT1: Completed match for expect: CONNECT

CHAT1: Chat script callout finished, status = Success
```

Chat script problems can be broken into three categories:

Configuration error

Modem failure

Connection failure

Chat Script Failure

This list shows possible outputs from debug chat shows and suggested actions:

```
no matching chat script found for [number]
```

A chat script has not been configured. Add one.

```
Chat script dialout finished, status = Connection timed out; remote host not responding
```

The modem is not responding to the chat script. Verify communication with the modem (refer to Table 16-2 in Chapter 16).

```
Timeout expecting: CONNECT
```

Possibility 1: The local modem is not actually placing the call. Verify that the modem can place a call by performing a reverse Telnet to the modem and manually initiating a dial.

Possibility 2: The remote modem is not answering. Test this by dialing the remote modem with an ordinary POTS telephone.

Possibility 3: The number being dialed is incorrect. Verify the number by dialing it manually. Correct the configuration, if necessary.

Possibility 4: The modem trainup is taking too long or the TIMEOUT value is too low. If the local modem is external, turn up the modem speaker volume and listen to the trainup tones. If the trainup is abruptly cut off, try increasing the TIMEOUT value in the chat-script command. If the TIMEOUT is already 60 seconds or more, see the Modem Trainup section.

ISDN Outbound Calling

Upon the first suspicion of an ISDN failure, either on a BRI or a PRI, always check the output from show isdn status. The key things to note are that Layer 1 should be Active and Layer 2 should be in a state of MULTIPLE_FRAME_ESTABLISHED. See the "Interpreting Show ISDN Status Output" section in Chapter 16 for information on reading this output, as well as for corrective measures.

For outbound ISDN calls, debug isdn q931 and debug isdn events are the best tools to use. Fortunately, debugging outbound calls is very similar to debugging incoming calls. A normal successful call might look like this:

```
*Mar 20 21:07:45.025: ISDN BR0: Event: Call to 5553759 at 64 Kb/s

*Mar 20 21:07:45.033: ISDN BR0: TX ->  SETUP pd = 8  callref = 0x2C

*Mar 20 21:07:45.037:        Bearer Capability i = 0x8890

*Mar 20 21:07:45.041:        Channel ID i = 0x83

*Mar 20 21:07:45.041:        Keypad Facility i = 0x35353533373539

*Mar 20 21:07:45.141: ISDN BR0: RX <-  CALL_PROC pd = 8  callref = 0xAC
```

```
*Mar 20 21:07:45.145:        Channel ID i = 0x89

*Mar 20 21:07:45.157: ISDN BR0: received HOST_PROCEEDING

       Channel ID i = 0x0101

*Mar 20 21:07:45.161:   -------------------

       Channel ID i = 0x89

*Mar 20 21:07:45.313: ISDN BR0: RX <-  CONNECT pd = 8  callref = 0xAC

*Mar 20 21:07:45.325: ISDN BR0: received HOST_CONNECT


!--- The CONNECT message is the key indicator of success. If a CONNECT is not received,

!--- you may see a DISCONNECT or a RELEASE_COMP (release complete) message followed by

!--- a cause code (see below)


*Mar 20 22:11:03.212: ISDN BR0: RX <-  RELEASE_COMP pd = 8  callref = 0x8F

*Mar 20 22:11:03.216:        Cause i = 0x8295 - Call rejected
```

The cause value indicates two things.

The second byte of the 4- or 6-byte value indicates from where in the end-to-end call path the DISCONNECT or RELEASE_COMP was received. This can help you to localize the problem.

The third and fourth bytes indicate the actual reason for the failure. See the tables which follow for the meanings of the different values.

Note: The following printout usually indicates a higher-protocol failure:

```
Cause i = 0x8090 - Normal call clearing
```

PPP authentication failure is a typical reason. Turn on debug ppp negotiation and debug ppp authentication before assuming that the connection failure is necessarily an ISDN problem

Cause Code Fields

Table 17-9 lists the ISDN cause code fields that display in the following format within the debug commands:

```
i=0x y1 y2 z1 z2 [a1 a2]
```

ISDN Cause Code Fields

Field	Value Description
0x	The values that follow are in hexadecimal.
y1	8--ITU-T standard coding.
y2	0--User 1--Private network serving local user 2--Public network serving local user 3--Transit network 4--Public network serving remote user 5--Private network serving remote user 7--International network A--Network beyond internetworking point
z1	Class (the more significant hexadecimal number) of cause value. Refer to the next table for detailed information about possible values.
z2	Value (the less significant hexadecimal number) of cause value. Refer to the next table for detailed information about possible values.
a1	(Optional) Diagnostic field that is always 8.
a2	(Optional) Diagnostic field that is one of the following values: 0--Unknown 1--Permanent 2--Transient

ISDN Cause Values

The following table lists descriptions of some of the most commonly-seen cause values of the cause information element - the third and fourth bytes of the cause code. For more complete information about ISDN codes and values, refer to Understanding debug isdn q931 Disconnect Cause Codes.

Hex Value	Cause	Explanation
81	Unallocated (unassigned) number	The ISDN number was sent to the switch in the correct format; however, the number is not assigned to any destination equipment.
90	Normal call clearing	Normal call clearing has occurred.
91	User busy	The called system acknowledges the connection request but is unable to accept the call because all B channels are in use.
92	No user responding	The connection cannot be completed because the destination does not respond to the call.
93	No answer from user (user alerted)	The destination responds to the connection request but fails to complete the connection within the prescribed time. The problem is at the remote end of the connection.
95	Call rejected	The destination is capable of accepting the call but rejected it for an unknown reason.
9C	Invalid number format	The connection could be not established because the destination address was presented in an unrecognizable format or because the destination address was incomplete.
9F	Normal,	Reports the occurrence of a normal event when no standard

	unspecified	cause applies. No action required.
A2	No circuit/channel available	The connection cannot be established because no appropriate channel is available to take the call.
A6	Network out of order	The destination cannot be reached because the network is not functioning correctly, and the condition might last for an extended period of time. An immediate reconnect attempt will probably be unsuccessful.
AC	Requested circuit/channel not available	The remote equipment cannot provide the requested channel for an unknown reason. This might be a temporary problem.
B2	Requested facility not subscribed	The remote equipment supports the requested supplementary service by subscription only. This is frequently a reference to long-distance service.
B9	Bearer capability not authorized	The user requested a bearer capability that the network provides, but the user is not authorized to use it. This might be a subscription problem.
D8	Incompatible destination	Indicates that an attempt was made to connect to non-ISDN equipment. For example, to an analog line.
E0	Mandatory information element is missing	The receiving equipment received a message that did not include one of the mandatory information elements. This is usually due to a D-channel error.

		If this error occurs systematically, report it to your ISDN service provider.
E4	Invalid information element contents	The remote equipment received a message that includes invalid information in the information element. This is usually due to a D-channel error.

CAS Outbound Calling

For outbound calling via CAS T1 or E1 and integrated digital modems, much of the troubleshooting is similar to other DDR troubleshooting. The same holds true, as well, for outbound integrated modem calls over a PRI line. The unique features involved in making a call in this manner require special debugging in the event of a call failure.

As for other DDR situations, you must ensure that a call attempt is demanded. Use debug dialer events for this purpose. Refer to Verifying Dialer Operation.

Before a call can be placed, a modem must be allocated for the call. To view this process, and the subsequent call, use the following debug commands:

```
debug modem

debug modem csm

debug cas
```

Note: The debug cas command first appeared in IOS version 12.0(7)T for the AS5200 and AS5300. Earlier versions of IOS use a system-level configuration command service internal along with the exec command modem-mgmt debug rbs:

Turning on the Debugs

```
router#conf t


Enter configuration commands, one per line.  End with CNTL/Z.

router(config)#service internal

router(config)#^Z


router#modem-mgmt csm ?
```

```
  debug-rbs     enable rbs debugging
  no-debug-rbs  disable rbs debugging
```

```
router#modem-mgmt csm debug-rbs
router#
neat msg at slot 0: debug-rbs is on
neat msg at slot 0: special debug-rbs is on
```

Turning off the Debugs

```
router#
router#modem-mgmt csm no-debug-rbs
neat msg at slot 0: debug-rbs is off
```

Note: Debugging this information on an AS5800 requires connecting to the trunk card. The following is an example of a normal outbound call over a CAS T1 that is provisioned and configured for FXS-Ground-Start:

```
Mica Modem(1/0): Rcvd Dial String(5551111) [Modem receives digits from chat script]
CSM_PROC_IDLE: CSM_EVENT_MODEM_OFFHOOK at slot 1, port 0
CSM_RX_CAS_EVENT_FROM_NEAT:(A003):  EVENT_CHANNEL_LOCK at slot 1 and port 0
CSM_PROC_OC4_DIALING: CSM_EVENT_DSX0_BCHAN_ASSIGNED at slot 1, port 0
Mica Modem(1/0): Configure(0x1)
Mica Modem(1/0): Configure(0x2)
Mica Modem(1/0): Configure(0x5)
Mica Modem(1/0): Call Setup
neat msg at slot 0: (0/2): Tx RING_GROUND
Mica Modem(1/0): State Transition to Call Setup
neat msg at slot 0: (0/2): Rx TIP_GROUND_NORING [Telco switch goes OFFHOOK]
CSM_RX_CAS_EVENT_FROM_NEAT:(A003):  EVENT_START_TX_TONE at slot 1 and port 0
CSM_PROC_OC5_WAIT_FOR_CARRIER: CSM_EVENT_DSX0_START_TX_TONE at slot 1, port 0
neat msg at slot 0: (0/2): Tx LOOP_CLOSURE [Now the router goes OFFHOOK]
```

```
Mica Modem(1/0): Rcvd Tone detected(2)

Mica Modem(1/0): Generate digits:called_party_num=5551111 len=8

Mica Modem(1/0): Rcvd Digits Generated

CSM_PROC_OC5_WAIT_FOR_CARRIER: CSM_EVENT_ADDR_INFO_COLLECTED at slot 1, port 0

CSM_RX_CAS_EVENT_FROM_NEAT:(A003):  EVENT_CHANNEL_CONNECTED at slot 1 and port 0

CSM_PROC_OC5_WAIT_FOR_CARRIER: CSM_EVENT_DSX0_CONNECTED at slot 1, port 0

Mica Modem(1/0): Link Initiate

Mica Modem(1/0): State Transition to Connect

Mica Modem(1/0): State Transition to Link

Mica Modem(1/0): State Transition to Trainup

Mica Modem(1/0): State Transition to EC Negotiating

Mica Modem(1/0): State Transition to Steady State

Mica Modem(1/0): State Transition to Steady State Speedshifting

Mica Modem(1/0): State Transition to Steady State
```

Debugs for T1s and E1s with other signaling types are similar.

Getting to this point in the debugging indicates that the calling and answering modems have trained and connected, and that higher-layer protocols can begin to negotiate. If a modem is properly allocated for the outbound call but the connection fails to get this far, the T1 must be examined. Refer to Chapter 15 for T1 troubleshooting information.

Troubleshooting PPP

Troubleshooting the PPP portion of a connection begins when you know that the dial connection, ISDN or async, successfully establishes.

It is important to understand what a successful debug PPP sequence looks like before you troubleshoot PPP negotiation. In this way, comparing a faulty PPP debug session against a successfully-completed debug PPP sequence saves you time and effort.

Following is an example of a successful PPP sequence. See PPP LCP Negotiation Details for a detailed description of the output fields.

```
Montecito#

Mar 13 10:57:13.415: %LINK-3-UPDOWN: Interface Async1, changed state to up

Mar 13 10:57:15.415: As1 LCP: O CONFREQ [ACKrcvd] id 2 len 25
Mar 13 10:57:15.415: As1 LCP:    ACCM 0x000A0000 (0x0206000A0000)

Mar 13 10:57:15.415: As1 LCP:    AuthProto CHAP (0x0305C22305)

Mar 13 10:57:15.415: As1 LCP:    MagicNumber 0x1084F0A2 (0x05061084F0A2)

Mar 13 10:57:15.415: As1 LCP:    PFC (0x0702)

Mar 13 10:57:15.415: As1 LCP:    ACFC (0x0802)

Mar 13 10:57:15.543: As1 LCP: I CONFACK [REQsent] id 2 len 25

Mar 13 10:57:15.543: As1 LCP:    ACCM 0x000A0000 (0x0206000A0000)

Mar 13 10:57:15.543: As1 LCP:    AuthProto CHAP (0x0305C22305)

Mar 13 10:57:15.543: As1 LCP:    MagicNumber 0x1084F0A2 (0x05061084F0A2)

Mar 13 10:57:15.543: As1 LCP:    PFC (0x0702)

Mar 13 10:57:15.547: As1 LCP:    ACFC (0x0802)

Mar 13 10:57:16.919: As1 LCP: I CONFREQ [ACKrcvd] id 4 len 23

Mar 13 10:57:16.919: As1 LCP:    ACCM 0x000A0000 (0x0206000A0000)

Mar 13 10:57:16.919: As1 LCP:    MagicNumber 0x001327B0 (0x0506001327B0)

Mar 13 10:57:16.919: As1 LCP:    PFC (0x0702)

Mar 13 10:57:16.919: As1 LCP:    ACFC (0x0802)

Mar 13 10:57:16.919: As1 LCP:    Callback 6  (0x0D0306)

Mar 13 10:57:16.919: As1 LCP: O CONFREJ [ACKrcvd] id 4 len 7

Mar 13 10:57:16.919: As1 LCP:    Callback 6  (0x0D0306)

Mar 13 10:57:17.047: As1 LCP: I CONFREQ [ACKrcvd] id 5 len 20

Mar 13 10:57:17.047: As1 LCP:    ACCM 0x000A0000 (0x0206000A0000)

Mar 13 10:57:17.047: As1 LCP:    MagicNumber 0x001327B0 (0x0506001327B0)

Mar 13 10:57:17.047: As1 LCP:    PFC (0x0702)

Mar 13 10:57:17.047: As1 LCP:    ACFC (0x0802)

Mar 13 10:57:17.047: As1 LCP: O CONFACK [ACKrcvd] id 5 len 20

Mar 13 10:57:17.047: As1 LCP:    ACCM 0x000A0000 (0x0206000A0000)

Mar 13 10:57:17.047: As1 LCP:    MagicNumber 0x001327B0 (0x0506001327B0)
Mar 13 10:57:17.047: As1 LCP:    PFC (0x0702)

Mar 13 10:57:17.047: As1 LCP:    ACFC (0x0802)

Mar 13 10:57:17.047: As1 LCP: State is Open

Mar 13 10:57:17.047: As1 PPP: Phase is AUTHENTICATING, by this end

Mar 13 10:57:17.047: As1 CHAP: O CHALLENGE id 1 len 28 from "Montecito"

Mar 13 10:57:17.191: As1 CHAP: I RESPONSE id 1 len 30 from "Goleta"

Mar 13 10:57:17.191: As1 CHAP: O SUCCESS id 1 len 4

Mar 13 10:57:17.191: As1 PPP: Phase is UP

Mar 13 10:57:17.191: As1 IPCP: O CONFREQ [Closed] id 1 len 10

Mar 13 10:57:17.191: As1 IPCP:    Address 172.22.66.23 (0x0306AC164217)

Mar 13 10:57:17.303: As1 IPCP: I CONFREQ [REQsent] id 1 len 40

Mar 13 10:57:17.303: As1 IPCP:    CompressType VJ 15 slots CompressSlotID

 (0x0206002D0F01)

Mar 13 10:57:17.303: As1 IPCP:    Address 0.0.0.0 (0x030600000000)

Mar 13 10:57:17.303: As1 IPCP:    PrimaryDNS 0.0.0.0 (0x810600000000)

Mar 13 10:57:17.303: As1 IPCP:    PrimaryWINS 0.0.0.0 (0x820600000000)

Mar 13 10:57:17.303: As1 IPCP:    SecondaryDNS 0.0.0.0 (0x830600000000)

Mar 13 10:57:17.303: As1 IPCP:    SecondaryWINS 0.0.0.0 (0x840600000000)

Mar 13 10:57:17.303: As1 IPCP: O CONFREJ [REQsent] id 1 len 22

Mar 13 10:57:17.303: As1 IPCP:    CompressType VJ 15 slots CompressSlotID

 (0x0206002D0F01)

Mar 13 10:57:17.303: As1 IPCP:    PrimaryWINS 0.0.0.0 (0x820600000000)

Mar 13 10:57:17.303: As1 IPCP:    SecondaryWINS 0.0.0.0 (0x840600000000)

Mar 13 10:57:17.319: As1 CCP: I CONFREQ [Not negotiated] id 1 len 15

Mar 13 10:57:17.319: As1 CCP:    MS-PPC supported bits 0x00000001 (0x120600000001)

Mar 13 10:57:17.319: As1 CCP:    Stacker history 1 check mode EXTENDED (0x1105000104)

Mar 13 10:57:17.319: As1 LCP: O PROTREJ [Open] id 3 len 21 protocol CCP
Mar 13 10:57:17.319: As1 LCP:  (0x80FD0101000F12060000000111050001)

Mar 13 10:57:17.319: As1 LCP:  (0x04)

Mar 13 10:57:17.319: As1 IPCP: I CONFACK [REQsent] id 1 len 10

Mar 13 10:57:17.319: As1 IPCP:    Address 172.22.66.23 (0x0306AC164217)

Mar 13 10:57:18.191: %LINEPROTO-5-UPDOWN: Line protocol on Interface Async1,

 changed state to up

Mar 13 10:57:19.191: As1 IPCP: TIMEout: State ACKrcvd

Mar 13 10:57:19.191: As1 IPCP: O CONFREQ [ACKrcvd] id 2 len 10

Mar 13 10:57:19.191: As1 IPCP:    Address 172.22.66.23 (0x0306AC164217)

Mar 13 10:57:19.315: As1 IPCP: I CONFACK [REQsent] id 2 len 10

Mar 13 10:57:19.315: As1 IPCP:    Address 172.22.66.23 (0x0306AC164217)

Mar 13 10:57:20.307: As1 IPCP: I CONFREQ [ACKrcvd] id 2 len 34

Mar 13 10:57:20.307: As1 IPCP:    Address 0.0.0.0 (0x030600000000)

Mar 13 10:57:20.307: As1 IPCP:    PrimaryDNS 0.0.0.0 (0x810600000000)

Mar 13 10:57:20.307: As1 IPCP:    PrimaryWINS 0.0.0.0 (0x820600000000)

Mar 13 10:57:20.307: As1 IPCP:    SecondaryDNS 0.0.0.0 (0x830600000000)

Mar 13 10:57:20.307: As1 IPCP:    SecondaryWINS 0.0.0.0 (0x840600000000)

Mar 13 10:57:20.307: As1 IPCP: O CONFREJ [ACKrcvd] id 2 len 16

Mar 13 10:57:20.307: As1 IPCP:    PrimaryWINS 0.0.0.0 (0x820600000000)

Mar 13 10:57:20.307: As1 IPCP:    SecondaryWINS 0.0.0.0 (0x840600000000)

Mar 13 10:57:20.419: As1 IPCP: I CONFREQ [ACKrcvd] id 3 len 22

Mar 13 10:57:20.419: As1 IPCP:    Address 0.0.0.0 (0x030600000000)

Mar 13 10:57:20.419: As1 IPCP:    PrimaryDNS 0.0.0.0 (0x810600000000)

Mar 13 10:57:20.419: As1 IPCP:    SecondaryDNS 0.0.0.0 (0x830600000000)

Mar 13 10:57:20.419: As1 IPCP: O CONFNAK [ACKrcvd] id 3 len 22

Mar 13 10:57:20.419: As1 IPCP:    Address 10.1.1.1 (0x03060A010101)

Mar 13 10:57:20.419: As1 IPCP:    PrimaryDNS 171.68.10.70 (0x8106AB440A46)
Mar 13 10:57:20.419: As1 IPCP:    SecondaryDNS 171.68.10.140 (0x8306AB440A8C)

Mar 13 10:57:20.543: As1 IPCP: I CONFREQ [ACKrcvd] id 4 len 22

Mar 13 10:57:20.543: As1 IPCP:    Address 10.1.1.1 (0x03060A010101)

Mar 13 10:57:20.547: As1 IPCP:    PrimaryDNS 171.68.10.70 (0x8106AB440A46)

Mar 13 10:57:20.547: As1 IPCP:    SecondaryDNS 171.68.10.140 (0x8306AB440A8C)

Mar 13 10:57:20.547: As1 IPCP: O CONFACK [ACKrcvd] id 4 len 22

Mar 13 10:57:20.547: As1 IPCP:    Address 10.1.1.1 (0x03060A010101)

Mar 13 10:57:20.547: As1 IPCP:    PrimaryDNS 171.68.10.70 (0x8106AB440A46)

Mar 13 10:57:20.547: As1 IPCP:    SecondaryDNS 171.68.10.140 (0x8306AB440A8C)

Mar 13 10:57:20.547: As1 IPCP: State is Open

Mar 13 10:57:20.551: As1 IPCP: Install route to 10.1.1.1
```

Note: Your debugs may appear in a different format. This example shows the newer PPP debugging output format which was modified in IOS version 11.2(8). See Chapter 16 for an example of PPP debugging with the older versions of IOS.

PPP LCP Negotiation Details

Time Stamp	Description
10:57:15.415	Outgoing configuration request (O CONFREQ). The NAS sends an outgoing PPP configuration request packet to the client.
10:57:15.543	Incoming configuration acknowledgment (I CONFACK). The client acknowledges Montecito's PPP request.
10:57:16.919	Incoming configuration request (I CONFREQ). The client wants to negotiate the callback protocol.
10:57:16.919	Outgoing configuration reject (O CONFREJ). The NAS rejects the callback option.

10:57:17.047	Incoming configuration request (I CONFREQ). The client requests a new set of options. Notice that Microsoft Callback is not requested this time.
10:57:17.047	Outgoing configuration acknowledgment (O CONFACK). The NAS accepts the new set of options.
10:57:17.047	PPP LCP negotiation is completed successfully. The LCP state is "Open". Both sides have acknowledged (CONFACK) the other side's configuration request (CONFREQ).
10:57:17.047 until 10:57:17.191	PPP authentication is completed successfully. After the LCP negotiates, authentication starts. Authentication must take place before any network protocols, such as IP, are delivered. Both sides authenticate with the method negotiated during LCP. Montecito is authenticating the client using CHAP.
10:57:20.551	The state is open for IP Control Protocol (IPCP). A route is negotiated and installed for the IPCP peer, which is assigned IP address 1.1.1.1.

Link Control Protocol

Two types of problems are typically encountered during LCP negotiation.

The first occurs when one peer makes configuration requests which the other peer cannot or will not acknowledge. While this is a frequent occurrence, it can be a problem if the requester insists on the parameter. A typical example is when negotiating AUTHTYPE (also known as "AuthProto"). For instance, many access servers are configured to accept only CHAP for authentication. If the caller is configured to do only PAP authentication, CONFREQs and CONFNAKs will be exchanged until one peer or the other drops the connection.

```
BR0:1 LCP: I CONFREQ [ACKrcvd] id 66 len 14

BR0:1 LCP:    AuthProto PAP (0x0304C023)
```

```
BR0:1 LCP:    MagicNumber 0xBC6B9F91 (0x0506BC6B9F91)

BR0:1 LCP: O CONFNAK [ACKrcvd] id 66 len 9

BR0:1 LCP:    AuthProto CHAP (0x0305C22305)

BR0:1 LCP: I CONFREQ [ACKrcvd] id 67 len 14

BR0:1 LCP:    AuthProto PAP (0x0304C023)

BR0:1 LCP:    MagicNumber 0xBC6B9F91 (0x0506BC6B9F91)

BR0:1 LCP: O CONFNAK [ACKrcvd] id 67 len 9

BR0:1 LCP:    AuthProto CHAP (0x0305C22305)

BR0:1 LCP: I CONFREQ [ACKrcvd] id 68 len 14

BR0:1 LCP:    AuthProto PAP (0x0304C023)

BR0:1 LCP:    MagicNumber 0xBC6B9F91 (0x0506BC6B9F91)

BR0:1 LCP: O CONFNAK [ACKrcvd] id 68 len 9

BR0:1 LCP:    AuthProto CHAP (0x0305C22305)

...

...
```

The second type of problem in LCP is when only outbound CONFREQs are seen on one or both peers as in the example below. This is usually the result of what is referred to as a speed mismatch at the lower layer. This condition can occur in either async or ISDN DDR.

```
Jun 10 19:57:59.768: As5 PPP: Phase is ESTABLISHING, Active Open

Jun 10 19:57:59.768: As5 LCP: O CONFREQ [Closed] id 64 len 25

Jun 10 19:57:59.768: As5 LCP: ACCM 0x000A0000 (0x0206000A0000)

Jun 10 19:57:59.768: As5 LCP: AuthProto CHAP (0x0305C22305)

Jun 10 19:57:59.768: As5 LCP: MagicNumber 0x5779D9D2 (0x05065779D9D2)

Jun 10 19:57:59.768: As5 LCP: PFC (0x0702)

Jun 10 19:57:59.768: As5 LCP: ACFC (0x0802)

Jun 10 19:58:01.768: As5 LCP: TIMEout: State REQsent

Jun 10 19:58:01.768: As5 LCP: O CONFREQ [REQsent] id 65 len 25

Jun 10 19:58:01.768: As5 LCP: ACCM 0x000A0000 (0x0206000A0000)
Jun 10 19:58:01.768: As5 LCP: AuthProto CHAP (0x0305C22305)
Jun 10 19:58:01.768: As5 LCP: MagicNumber 0x5779D9D2 (0x05065779D9D2)
Jun 10 19:58:01.768: As5 LCP: PFC (0x0702)
Jun 10 19:58:01.768: As5 LCP: ACFC (0x0802).
Jun 10 19:58:03.768: As5 LCP: TIMEout: State REQsent
Jun 10 19:58:03.768: As5 LCP: O CONFREQ [REQsent] id 66 len 25
Jun 10 19:58:03.768: As5 LCP: ACCM 0x000A0000 (0x0206000A0000)
Jun 10 19:58:03.768: As5 LCP: AuthProto CHAP (0x0305C22305)
Jun 10 19:58:03.768: As5 LCP: MagicNumber 0x5779D9D2 (0x05065779D9D2)
Jun 10 19:58:03.768: As5 LCP: PFC (0x0702)
Jun 10 19:58:03.768: As5 LCP: ACF.C (0x0802)
Jun 10 19:58:05.768: As5 LCP: TIMEout: State REQsent
Jun 10 19:58:05.768: As5 LCP: O CONFREQ [REQsent] id 67 len 25


!--- This repeats every two seconds until:


Jun 10 19:58:19.768: As5 LCP: O CONFREQ [REQsent] id 74 len 25
Jun 10 19:58:19.768: As5 LCP: ACCM 0x000A0000 (0x0206000A0000)
Jun 10 19:58:19.768: As5 LCP: AuthProto CHAP (0x0305C22305)
Jun 10 19:58:19.768: As5 LCP: MagicNumber 0x5779D9D2 (0x05065779D9D2)
Jun 10 19:58:19.768: As5 LCP: PFC (0x0702)
Jun 10 19:58:19.768: As5 LCP: ACFC (0x0802)
Jun 10 19:58:21.768: As5 LCP: TIMEout: State REQsent
Jun 10 19:58:21.768: TTY5: Async Int reset: Dropping DTR
```

If the connection is async, the probable cause is a speed mismatch between the router and its modem. This is usually as a result of having failed to lock the DTE speed of the modem to the configured speed of the TTY line. The problem may be found on either or both of the peers, so check both. Refer to Modem Cannot Send or Receive Data earlier in this chapter.

If the symptoms are seen when the connection is over ISDN, the problem is likely to be that one peer is connecting at 56K while the other is at 64K. While this condition is rare, it does happen. The problem could be one or both peers, or possibly the telephone company. Use debug isdn q931 and examine the SETUP messages on each of the peers. The Bearer Capability sent from one peer should match the Bearer Capability seen in the SETUP message received on the other peer. As a possible remedy, configure the dialing speed, 56K or 64K, in either the interface level command dialer map or in the command dialer isdn speed configured under a map-class.

```
*Mar 20 21:07:45.033: ISDN BR0: TX ->  SETUP pd = 8  callref = 0x2C

*Mar 20 21:07:45.037:        Bearer Capability i = 0x8890

*Mar 20 21:07:45.041:        Channel ID i = 0x83

*Mar 20 21:07:45.041:        Keypad Facility i = 0x35353533373539
```

This situation is one which may warrant a call to the Cisco TAC. Collect the following outputs from both peers before calling the TAC:

```
show running-config

show version

debug isdn q931

debug isdn events

debug ppp negotiation
```

Authentication

Failed authentication is the single most common reason for a PPP failure. Misconfigured or mismatched usernames and passwords create error messages in debug output.

The following example shows that the username Goleta does not have permission to dial in to the NAS, which does not have a local username configured for this user. To fix the problem, use the username name password password command to add the username "Goleta" to the NAS' local AAA database:

```
Mar 13 11:01:42.399: As2 LCP: State is Open

Mar 13 11:01:42.399: As2 PPP: Phase is AUTHENTICATING, by this end

Mar 13 11:01:42.399: As2 CHAP: O CHALLENGE id 1 len 28 from "Montecito"

Mar 13 11:01:42.539: As2 CHAP: I RESPONSE id 1 len 30 from "Goleta"

Mar 13 11:01:42.539: As2 CHAP: Unable to validate Response.  Username Goleta not found

Mar 13 11:01:42.539: As2 CHAP: O FAILURE id 1 len 26 msg is "Authentication failure"
Mar 13 11:01:42.539: As2 PPP: Phase is TERMINATING
```

The following example shows that the username "Goleta" is configured on the NAS. However, the password comparison failed. To fix this problem, use the username name password password command to specify the correct login password for Goleta:

```
Mar 13 11:04:06.843: As3 LCP: State is Open

Mar 13 11:04:06.843: As3 PPP: Phase is AUTHENTICATING, by this end

Mar 13 11:04:06.843: As3 CHAP: O CHALLENGE id 1 len 28 from "Montecito"

Mar 13 11:04:06.987: As3 CHAP: I RESPONSE id 1 len 30 from "Goleta"

Mar 13 11:04:06.987: As3 CHAP: O FAILURE id 1 len 25 msg is "MD/DES compare failed"

Mar 13 11:04:06.987: As3 PPP: Phase is TERMINATING
```

For more information on PAP authentication refer to Configuring and Troubleshooting PPP Password Authentication Protocol (PAP).

Network Control Protocol

After the peers have successfully performed the required authentication, the negotiation moves into the NCP phase. If both peers are properly configured, the NCP negotiation might look like the following example which shows a client PC dialing into and negotiating with a NAS:

```
solvang# show debug

Generic IP:

IP peer address activity debugging is on

PPP:

PPP protocol negotiation debugging is on



*Mar  1 21:35:04.186: As4 PPP: Phase is UP

*Mar  1 21:35:04.190: As4 IPCP: O CONFREQ [Not negotiated] id 1 len 10

*Mar  1 21:35:04.194: As4 IPCP:    Address 10.1.2.1 (0x03060A010201)

*Mar  1 21:35:04.282: As4 IPCP: I CONFREQ [REQsent] id 1 len 28

*Mar  1 21:35:04.282: As4 IPCP:    CompressType VJ 15 slots CompressSlotID

 (0x0206002D0F01)

*Mar  1 21:35:04.286: As4 IPCP:    Address 0.0.0.0 (0x030600000000)
```

```
*Mar  1 21:35:04.290: As4 IPCP:    PrimaryDNS 0.0.0.0 (0x810600000000)
*Mar  1 21:35:04.298: As4 IPCP:    SecondaryDNS 0.0.0.0 (0x830600000000)
*Mar  1 21:35:04.306: As4 IPCP: O CONFREJ [REQsent] id 1 len 10
*Mar  1 21:35:04.310: As4 IPCP:    CompressType VJ 15 slots CompressSlotID
 (0x0206002D0F01)
*Mar  1 21:35:04.314: As4 CCP: I CONFREQ [Not negotiated] id 1 len 15
*Mar  1 21:35:04.318: As4 CCP:    MS-PPC supported bits 0x00000001 (0x120600000001)
*Mar  1 21:35:04.318: As4 CCP:    Stacker history 1 check mode EXTENDED (0x1105000104)
*Mar  1 21:35:04.322: As4 LCP: O PROTREJ [Open] id 3 len 21 protocol CCP
*Mar  1 21:35:04.326: As4 LCP:  (0x80FD0101000F120600000001110500001)
*Mar  1 21:35:04.330: As4 LCP:  (0x04)
*Mar  1 21:35:04.334: As4 IPCP: I CONFACK [REQsent] id 1 len 10
*Mar  1 21:35:04.338: As4 IPCP:    Address 10.1.2.1 (0x03060A010201)
*Mar  1 21:35:05.186: %LINEPROTO-5-UPDOWN: Line protocol on Interface Async4,
 changed state to up
*Mar  1 21:35:07.274: As4 IPCP: I CONFREQ [ACKrcvd] id 2 len 22
*Mar  1 21:35:07.278: As4 IPCP:    Address 0.0.0.0 (0x030600000000)
*Mar  1 21:35:07.282: As4 IPCP:    PrimaryDNS 0.0.0.0 (0x810600000000)
*Mar  1 21:35:07.286: As4 IPCP:    SecondaryDNS 0.0.0.0 (0x830600000000)
*Mar  1 21:35:07.294: As4 IPCP: O CONFNAK [ACKrcvd] id 2 len 22
*Mar  1 21:35:07.298: As4 IPCP:    Address 10.1.2.2 (0x03060A010202)
*Mar  1 21:35:07.302: As4 IPCP:    PrimaryDNS 10.2.2.3 (0x81060A020203)
*Mar  1 21:35:07.310: As4 IPCP:    SecondaryDNS 10.2.3.1 (0x83060A020301)
*Mar  1 21:35:07.426: As4 IPCP: I CONFREQ [ACKrcvd] id 3 len 22
*Mar  1 21:35:07.430: As4 IPCP:    Address 10.1.2.2 (0x03060A010202)
*Mar  1 21:35:07.434: As4 IPCP:    PrimaryDNS 10.2.2.3 (0x81060A020203)
*Mar  1 21:35:07.442: As4 IPCP:    SecondaryDNS 10.2.3.1 (0x83060A020301)
```

```
*Mar  1 21:35:07.446: ip_get_pool: As4: validate address = 10.1.2.2

*Mar  1 21:35:07.450: ip_get_pool: As4: using pool default

*Mar  1 21:35:07.450: ip_get_pool: As4: returning address = 10.1.2.2

*Mar  1 21:35:07.454: set_ip_peer_addr: As4: address = 10.1.2.2 (3) is redundant

*Mar  1 21:35:07.458: As4 IPCP: O CONFACK [ACKrcvd] id 3 len 22

*Mar  1 21:35:07.462: As4 IPCP:    Address 10.1.2.2 (0x03060A010202)

*Mar  1 21:35:07.466: As4 IPCP:    PrimaryDNS 10.2.2.3 (0x81060A020203)

*Mar  1 21:35:07.474: As4 IPCP:    SecondaryDNS 10.2.3.1 (0x83060A020301)

*Mar  1 21:35:07.478: As4 IPCP: State is Open

*Mar  1 21:35:07.490: As4 IPCP: Install route to 10.1.2.2
```

PPP NCP Negotiation Details

Time Stamp	Description
21:35:04.190	Outgoing configuration request (O CONFREQ). The NAS sends an outgoing PPP configuration request packet containing its IP address to the peer.
21:35:04.282	Incoming CONFREQ. The peer requests to do VJ header compression. It needs an IP address for itself, as well as addresses of the primary and secondary DNS servers.
21:35:04.306	Outbound Config-Reject (CONFREJ). VJ header compression is rejected.
21:35:04.314 until 21:35:04.330	The peer sends a request to do Compression Control Protocol; the entire protocol is rejected by the NAS by means of a PROTREJ message. The peer should not (and does not) attempt to retry CCP.
21:35:04.334	The peer acknowledges the IP address of the NAS with a CONFACK.

21:35:07.274	Incoming CONFREQ. The peer no longer requests to do VJ header compression, but still needs an IP address for itself, as well as addresses of the primary and secondary DNS servers.
21:35:07.294	The NAS sends a CONFNAK containing the address it wants the peer to use, and addresses of the primary and secondary DNS servers.
21:35:07.426	The peer sends the addresses back to the NAS; an attempt to confirm that the addresses were properly received.
21:35:07.458	The NAS acknowledges the addresses with a CONFACK.
21:35:07.478	Each side of the connection having issued a CONFACK, negotiation finishes. The command show interfaces Async4 on the NAS shows "IPCP: Open".
21:35:07.490	A host route to the remote peer is installed in the NAS' routing table.

It is possible for the peers to simultaneously negotiate more than one Layer 3 protocol. It is not uncommon, for instance, to see IP and IPX being negotiated. It is also possible for one protocol to successfully negotiate while the other fails to do so.

Troubleshooting NCP

Any problems which occur during NCP negotiation can typically be traced to the configurations of the negotiating peers. If PPP negotiation fails during the NCP phase, refer to the following steps:

Verify interface protocol configuration

Examine the output of the privileged exec command show running-config. Verify that the interface is configured to support the protocol you wish to run over the connection.

Verify interface address

Confirm that the interface in question has an address configured. If using ip unnumbered [interface-name] or ipx ppp-client loopback [number], ensure that the referenced interface is configured with an address.

Verify client address availability

If the NAS is supposed to issue an IP address to the caller, ensure that such an address is available. The IP address to be handed out to the caller can be obtained through one of the following methods:

Configure locally on the interface. Check the interface configuration for the command peer default ip address a.b.c.d. In practice, this method should only be used on interfaces which accept connections from a single caller, such as on an async (not a group-async) interface.

Address pool locally configured on the NAS. The interface should have the command peer default ip address pool [pool-name]. In addition, the pool must be defined at the system level with the command ip local pool [pool-name] [first-address] [last-address]. The range of addresses defined in the pool should be large enough to accommodate as many simultaneously-connected callers as the NAS is capable.

DHCP server. The NAS interface must be configured with the command peer default ip address dhcp. Furthermore, the NAS must be configured to point to a DHCP server with the global configuration command ip dhcp-server [address].

AAA. If using TACACS+ or RADIUS for authorization, the AAA server can be configured to hand a specific IP address to a given caller every time that caller connects. See Chapter 16 for more information.

Verify server address configuration

To return the configured addresses of Domain Name Servers or Windows NT servers in response to BOOTP requests, ensure that the global-level commands async-bootp dns-server [address] and async-bootp nbns-server [address] are configured.

Note: While the command async-bootp subnet-mask [mask] can be configured on the NAS, the subnet mask will not be negotiated between the NAS and a PPP dial-in client PC. Due to the nature of point-to-point connections, the client automatically uses the IP address of the NAS (learned during IPCP negotiation) as the default gateway. The subnet mask is not needed in that point-to-point environment. The PC knows that if the destination address does not match the local address, the packet should be forwarded to the default gateway (NAS) which is always reached via the PPP link.

Before Calling the Cisco Systems TAC Team

Before calling the Cisco Systems Technical Assistance Center (TAC), make sure you have read through this chapter and completed the actions suggested for your system's problem.

Additionally, do the following and document the results so that we can better assist you:

For all problems, collect the output of show running-config and show version. Ensure that the command service timestamps debug datetime msec is in the configuration.

For DDR problems, collect the following:

show dialer map

debug dialer

debug ppp negotiation

debug ppp authentication

If ISDN is involved, collect:

show isdn status

debug isdn q931

debug isdn events

If modems are involved, collect:

show lines

show line [x]

show modem (if integrated modems are involved)

show modem version (if integrated modems are involved)

debug modem

debug modem csm (if integrated modems are involved)

debug chat (if a DDR scenario)

If T1s or PRIs are involved, collect:

show controller t1

What Is a TCP/IP Routing Table?

Each IP address identifies a remote router (or other *network gateway*) that the local router is configured to recognize. For each IP address, the routing table additionally stores a *network mask* and other data that specifies the destination IP address ranges that remote device will accept.

Home network routers utilize a very small routing table because they simply forward all outbound traffic to the Internet Service Provider (ISP) gateway which takes care of all other routing steps. Home router tables typically contain ten or fewer entries. By comparison, the largest routers at the core of the Internet backbone must maintain the full *Internet routing table* that exceeds 100,000 entries and growing as the Internet expands.

Two hypothetical, partial routing table entries are shown below:

IP Address: 172.48.11.181 - Network Mask: 255.255.255.255

IP Address: 192.168.1.1 - Network Mask: 255.255.255.0

In this example, the first entry represents the route to the ISP's primary DNS server. Requests made from the home network to any destination on the Internet will be sent to the IP address 172.48.11.181 for forwarding. The second entry represents the route between any computers within the home network, where the home router has IP address 192.168.1.1.

Dynamic vs. Static Routing

Home routers set up their routing tables automatically when connected to the ISP, a process called *dynamic routing*. They generate one routing table entry for each of the ISPs DNS servers (primary, secondary and tertiary if available) and one entry for routing among all the home computers. They may also generate a few additional routes for other special cases including *multicast* and *broadcast* routes.

Most residential network routers prevent you from manually overriding or changing the routing table. However, business routers typically allow network administrators to manually update or manipulate routing tables. This so-called *static routing* can be useful when optimizing for network performance and reliability.

Viewing the Contents of Routing Tables

On home broadband routers, the routing table contents are typically shown on a screen inside the administrative console.

On Windows and Unix/Linux computers, the *netstat -r* command also displays the contents of the routing table configured on the local computer.

Suggested Reading

- How to Set Up a Network Router
- What Is the Difference Between a Router and Hub / Switch?
- How to Find a Router's IP Address

Introduction to Demand-Dial Routing

Demand-dial routing is the forwarding of packets across a Point-to-Point Protocol (PPP) link. The PPP link is represented inside the Windows 2000 router as a demand-dial interface. Demand-dial interfaces can be used to create on-demand connections across dial-up, non-permanent or persistent media.

With local area network (LAN) and permanent wide area network (WAN) links, the interface that is being used to forward the packet is always in an active or connected state. The packet can be forwarded without having to create the physical or logical connection. However, the demand-dial interface can either be in a connected state or a disconnected state. If in a disconnected state when the packet is being forwarded, the demand-dial interface must be changed to a connected state before the packet can be forwarded.

The connection establishment process, consisting of creating a physical connection or a logical connection and a PPP connection, introduces a delay in the forwarding of the packet called the connection establishment delay. The length of the connection establishment delay varies for the type of physical or logical connection being established. For example, the connection establishment delay for analog phone lines or X.25 dialing in to a packet assembler-disassembler (PAD) can be 10 to 20 seconds or more. For Integrated Services Digital Network (ISDN) lines, the connection establishment delay can be as small as 3 to 5 seconds.

The connection establishment delay is an important consideration for applications being used across a demand-dial connection. There are two behaviors of applications to consider:

- How long it takes for the application to abandon the attempt to establish network communications, also known as application time-out. If the application time-out is longer than the connection establishment delay, then the application fails to establish communications and presents an error message to the user.
- How many times it attempts to establish network communications. On the first attempt, network traffic is forwarded to the demand-dial router which begins the connection establishment process. Due to the size of a finite buffer in the router, additional packets to be forwarded across the demand-dial connection that arrive during the connection establishment process might overwrite the initial application connection attempt packet. If the application tries to establish communications multiple times, then there is a better chance of forwarding an application connection attempt packet once the connection is established.

Applications that have long time-outs or multiple retries might not fail while waiting for the link to become available. Interactive applications such as Web browsers and Telnet might fail when first connecting. However, when the user retries the connection attempt, it succeeds because the first connection attempt created the demand-dial connection.

Once the connection is established, packets are forwarded across the demand-dial connection. Because the costs of demand-dial connections are typically time sensitive, after a configured amount of idle time the demand-dial link is terminated. Demand-dial connections have the benefit of allowing the user to use cheaper dial-up WAN links and only pay for the link when it is being used.

Demand-Dial Routing and Remote Access

Demand-dial routing is not the same as remote access; remote access connects a single user to a network, and demand-dial routing connects networks together. However, both remote access and demand-dial routing use PPP as the protocol mechanism to negotiate and authenticate the connection and encapsulate data sent on the connection. As implemented in the Windows 2000 Routing and Remote Access service, both remote access and demand-dial connections can be enabled separately but share the same:

- Behavior of the dial-in properties of user accounts.
- Security, including authentication protocols and encryption.
- Use of remote access policies.
- Use of Windows or Remote Authentication Dial-In User Service (RADIUS) as authentication providers.
- IP and Internetwork Packet Exchange (IPX) address allocation configuration.
- Use of PPP features such as Microsoft Point-to-Point Compression (MPPC), Multilink, and Bandwidth Allocation Protocol (BAP).
- Troubleshooting facilities including event logging, Windows or RADIUS authentication and accounting logging, and tracing.

Internetwork Routing

The following terms are essential to your understanding of routing:

End Systems . As defined by the International Standards Organization (ISO), end systems are network devices without the ability to forward packets between portions of a network. End systems are also known as hosts.

Intermediate Systems . Network devices with the ability to forward packets between portions of a network. Bridges, switches, and routers are examples of intermediate systems.

Network . A portion of the networking infrastructure (encompassing repeaters, hubs, and bridges/Layer 2 switches that is bound by a network layer intermediate system and is associated with the same network layer address.

Router . A network layer intermediate system used to connect networks together based on a common network layer protocol.

Hardware Router . A router that performs routing as a dedicated function and has specific hardware designed and optimized for routing.

Software Router . A router that is not dedicated to performing routing but performs routing as one of multiple processes running on the router computer. The Windows 2000 Server Router Service is a software router.

Internetwork . At least two networks connected using routers. Figure 1.1 illustrates an internetwork.

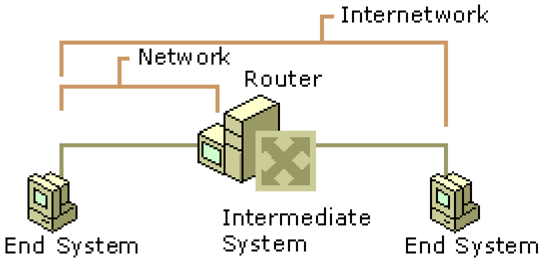

Figure 1.1 An Internetwork

Addressing in an Internetwork

The following internetwork addressing terms are also important to your understanding of routing:

Network address . Also known as a network ID. The number assigned to a single network in an internetwork. Network addresses are used by hosts and routers when routing a packet from a source to a destination in an internetwork.

Host address . Also known as a host ID or a node ID. Can either be the host's physical address (the address of the network interface card) or an administratively assigned address that uniquely identifies the host on its network.

Internetwork address . The combination of the network address and the host address; it uniquely identifies a host on an internetwork. An IP address that contains a network ID and a host ID is an internetwork address.

For detailed information about how IP implements network ID and host ID addressing, see "Introduction to TCP/IP" in the *Microsoft ® Windows ® 2000 Server Resource Kit TCP/IP Core Networking Guide* .

When a packet is sent from a source host to a destination host on an internetwork, the Network layer header of the packet contains:

- The Source Internetwork Address, which contains a source network address and source host address.
- The Destination Internetwork Address, which contains a destination network address and destination host address.
- A Hop Count, which either starts at zero and increases numerically for each router crossed to a maximum value, or starts at a maximum value and decreases numerically to zero for each router crossed. The hop count is used to prevent the packet from endlessly circulating on the internetwork.

- ## Routing Concepts
- Routing is the process of transferring data across an internetwork from a source host to a destination host. Routing can be understood in terms of two processes: host routing and router routing .
- Host routing occurs when the sending host forwards a packet. Based on the destination network address, the sending host must decide whether to forward the packet to the destination or to a router. In Figure 1.2, the Source Host forwards the packet destined for the Destination Host to Router 1.
- Router routing occurs when a router receives a packet that is to be forwarded. The packet is forwarded between routers (when the destination network is not directly attached to the router) or between a router and the destination host (when the destination network is directly attached). In Figure 1.2, Router 1 forwards the packet to Router 2. Router 2 forwards the packet to the Destination Host.
- 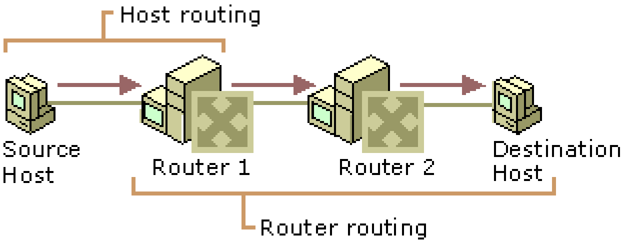

- **Figure 1.2 The Routing Process**

Foundations of Routing Protocols

Dynamic routers use routing protocols to facilitate the ongoing communication and dynamic updating of routing tables. Routing protocols are used between routers and represent additional network traffic overhead on the network. This additional traffic can become an important factor in planning WAN link usage. RIP and OSPF for IP, and RIP and NLSP for IPX are all routing protocols. In some cases, such as RIP for IP (version 1) and RIP for IPX, the routing information is exchanged using MAC-level broadcasts.

An important element of a routing protocol implementation is its ability to sense and recover from internetwork faults. How quickly it can recover is determined by the type of fault, how it is sensed, and how the routing information is propagated through the internetwork.

When all the routers on the internetwork have the correct routing information in their routing tables, the internetwork has converged. When convergence is achieved, the internetwork is in a stable state and all routing occurs along optimal paths.

When a link or router fails, the internetwork must reconfigure itself to reflect the new topology. Information in routing tables must be updated. Until the internetwork reconverges, it is in an unstable state in which routing loops and black holes can occur. The time it takes for the internetwork to reconverge is known as the convergence time. The convergence time varies based on the routing protocol and the type of failure (downed link or downed router).

Routing protocols are based either on a distance vector or link state technology. The main differences between distance vector and link state routing protocols include the following:

- What routing information is exchanged.
- How the information is exchanged.
- How quickly the internetwork can recover from a downed link or a downed router.

- ## Routing Infrastructure
- The routing infrastructure is the entire structure of the routed internetwork. The infrastructure has important attributes to consider when you are deciding on which routable protocols and routing protocols to use.
- ## Routers and Broadcast Traffic
- Internetwork-level broadcasts are Media Access Control (MAC)-level broadcast frames with a special destination internetwork address that informs the router that the packet is to be forwarded to all other networks except the network on which it was received. Routers must be configured to pass internetwork-level broadcast traffic. A MAC-level broadcast frame is used to reach all the hosts on a

network. Routers, unlike bridges, do not forward MAC-level broadcast traffic. However, to reach all the hosts on an internetwork, some routable protocols support the use of internetwork-level broadcasts.

- The inherent danger of forwarding internetwork-level broadcasts is the possibility of an internetwork-level broadcast storm in which a host malfunctions and continuously sends out the same internetwork-level broadcast packet. If the routers forward this traffic, the result is that all the hosts on the internetwork process each broadcast frame, possibly crippling the entire internetwork.
- The NetBIOS over IPX broadcast is an internetwork-level broadcast. NetBIOS applications on an IPX internetwork use a NetBIOS over IPX broadcast to perform name registration, resolution, and release. When the NetBIOS over IPX broadcast packet is received by an IPX router, the router records the network on which the packet was received in the NetBIOS over IPX header. Thus, the internetwork path is recorded in the NetBIOS over IPX header as it traverses the IPX internetwork.
- Before being forwarded, the IPX router checks the internetwork path information to prevent the forwarding of the NetBIOS over IPX broadcast onto a network on which it has already traveled. This prevents the broadcast from looping and causing more broadcast traffic. As an additional safeguard, NetBIOS over IPX broadcast packets can only propagate across eight networks using seven routers. At the eighth router, the packet is discarded without notifying the sending host. This is known as a silent discard. For more information about NetBIOS over IPX broadcasts, see "IPX Routing" in this book.
-
- **Note**
- An IPX internetwork path is recorded in a similar fashion to the MAC-sublayer routing information in a Token Ring source routing Explorer frame. However, unlike Token Ring source routing, the IPX internetwork path is not used in the subsequent communication. The IPX internetwork path is only used to prevent the broadcast packet from being forwarded on the same IPX network more than once

Tunneling

- Tunneling, also known as encapsulation, is a method of using an internetwork infrastructure of one protocol to transfer a payload. Typically, the payload is the frames (or packets) of another protocol (see Figure 1.8). Instead of being sent as it is produced by the originating host, the frame is encapsulated with an additional header. The additional header provides routing information so the encapsulated payload can traverse an intermediate internetwork (also known as a transit internetwork). The encapsulated packets are then routed between tunnel endpoints over the transit internetwork. Once the encapsulated payload packets reach their destination on the transit internetwork, the frame is de-encapsulated and forwarded to its final destination.
- The entire process of encapsulation, transmission, and de-encapsulation of packets is known as tunneling. The logical path through which the encapsulated packets travel through the transit internetwork is called a tunnel.
- 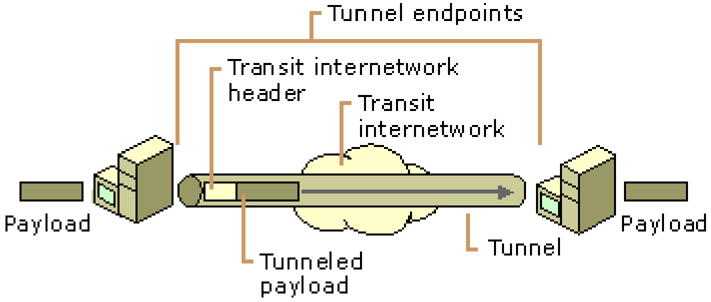

- **Figure 1.8 Tunneling**
- The transit internetwork can be any internetwork. The Internet is a good example as the most widely known public internetwork. There are also many examples of tunnels that are carried over corporate internetworks.
- Some common types of tunneling:
- **SNA Tunneling over IP Internetworks** To send System Network Architecture (SNA) traffic across a corporate IP internetwork, the SNA frame is encapsulated with a User Datagram Protocol (UDP) and IP header. This is known as Data Link Switching (DLSw) and is described in RFC 1795.
- **IPX Tunneling for Novell NetWare** IPX packets are sent to a NetWare server or IPX router that wraps the IPX packet with a UDP and IP header and sends them across an IP internetwork. The destination IP router removes the UDP and IP header and forwards them to the appropriate IPX destination.
- **Point-to-Point Tunneling Protocol** Point-to-Point Tunneling Protocol (PPTP) allows IP, IPX, or NetBEUI traffic to be encrypted and encapsulated in an IP header to be sent across a corporate IP internetwork or public internetworks like the Internet. For more information, see "Virtual Private Networking" in this book.
- **Layer 2 Tunneling Protocol** Layer Two Tunneling Protocol (L2TP) allows IP, IPX, or NetBEUI traffic to be encrypted and then sent over any medium that supports point-to-point datagram delivery such as IP, X.25, Frame Relay, or ATM. For more information, see "Virtual Private Networking" in this book.

- **IP Security (IPSec) Tunnel Mode** IPSec Tunnel Mode allows IP payloads to be encrypted and then encapsulated in an IP header to be sent across a corporate IP internetwork or public internetworks like the Internet. For more information about IPSec, see "Internet Protocol Security" in the *TCP/IP Core Networking Guide* .
- 📝

Router Routing

When a router is forwarded a packet that is not destined for that router, the router must either deliver it to the destination host or to another router, as shown in Figure 1.4.

- If the destination network matches a network to which the router is attached, the router forwards the packet to the destination host by addressing the packet to the destination host's physical address. The router performs a direct delivery to the destination.
- Conversely, if the destination network is not directly attached, the router forwards the packet to an intermediate router. The intermediate router chosen is based on the forwarding address of the optimal route in the routing table. The router forwards the packet by addressing the packet to the intermediate router's physical address. The router performs an indirect delivery to the next router in the path to the destination.

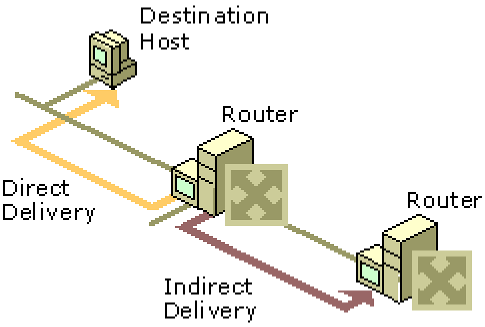

Figure 1.4 Router Routing Process

Top Of Page

Routing Problems

Routing problems can occur when either the host's or router's routing tables contain information that does not reflect the correct topology of the internetwork.

Routing Loops

During the router routing process, the packets are forwarded in the optimal direction according to the information in the local routing table. If the routing table entries on all the routers are correct, the packet takes the optimal path from the source to the destination. However, if any routing table entries are not correct, either through a misconfiguration or through learned routes that do not accurately reflect the topology of the internetwork, then routing loops can form. A routing loop is a path through the internetwork for a network ID that loops back onto itself.

Figure 1.6 illustrates a routing loop in which:

- According to the routing table on Router 1, the optimal route to Network 10 is through Router 2.
- According to the routing table on Router 2, the optimal route to Network 10 is through Router 3.
- According to the routing table on Router 3, the optimal route to Network 10 is through Router 1.

The hop count in the network layer header is used to prevent the packet from perpetually looping. Each time a router passes the packet from one network to another, it either increases or decreases the hop count. If the hop count reaches its maximum value (when increasing) or is 0 (when decreasing), the packet is discarded by the router.

For example, IPX hosts send IPX packets with a 0 hop count. Each RIP for IPX router increases the hop count by one. When it reaches 17, the packet is silently discarded. When IP hosts send IP packets, they set a maximum link count in the Time-to-Live (TTL) field in the IP header. Each IP router encountered decreases the TTL by one. When the TTL is 0, the IP router discards the packet and sends an ICMP Time Exceeded message back to the sending host. By default, Windows NT version 4.0 and later TCP/IP hosts set the TTL to the value of 128.

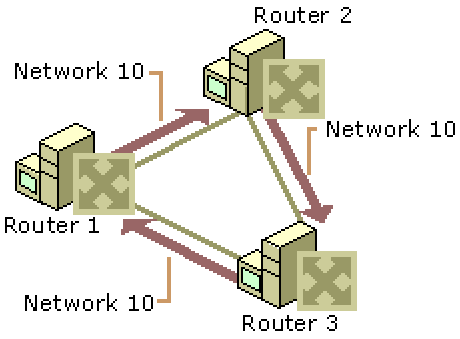

Figure 1.6 A Routing Loop

Top Of Page

Black Holes

Common internetworking protocols such as IP and IPX are connectionless, datagram-based protocols. They do not guarantee a successful delivery. IP and IPX attempt a best effort, unacknowledged delivery to the next hop or the final destination. This behavior can lead to conditions on the internetwork in which data is lost.

If a downstream router goes down and is not detected by the upstream router, the upstream router still forwards the packets to the downed router. Because the failed downstream router does not receive them, the packets forwarded by the upstream router are dropped from the internetwork. The upstream router is sending packets to a black hole, a condition of an internetwork where packets are lost without an indication of the error. In Figure 1.7, Router 1 has not been informed that Router 2 has failed and continues to forward packets to Router 2. The failed Router 2 creates a black hole.

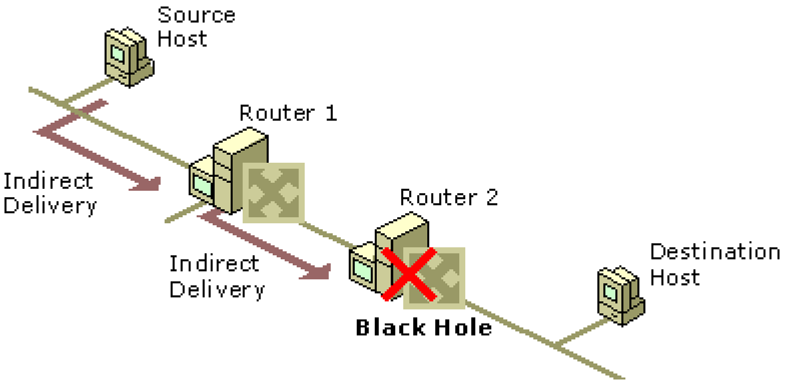

Figure 1.7 Routing Black Hole

Black holes can form when a link or router fails, and the failure is not yet detected. In a static routing environment, black holes persist until the link or router is brought back up or the static routers are reconfigured by the network administrator. In a dynamic routing environment, routers sense downed links or routers through the expiration of the lifetime of learned routes in their routing tables.

Black holes can also form when an active router discards packets without indicating the reason why the packets are being discarded. A good example is a Path Maximum Transmit Unit (PMTU) black hole router that discards IP packets that must be fragmented without sending a message to the sender indicating the error. PMTU black hole routers can be difficult to detect because packets of smaller sizes are forwarded. For more information about this specific issue, see the "TCP/IP Troubleshooting" chapter in the *TCP/IP Core Networking Guide* .

- ## Static and Dynamic Routers

- For routing between routers to work efficiently in an internetwork, routers must have knowledge of other network IDs or be configured with a default route. On large internetworks, the routing tables must be maintained so that the traffic always travels along optimal paths. How the routing tables are maintained defines the distinction between static and dynamic routing.
- **Static Routing**
- A router with manually configured routing tables is known as a static router. A network administrator, with knowledge of the internetwork topology, manually builds and updates the routing table, programming all routes in the routing table. Static routers can work well for small internetworks but do not scale well to large or dynamically changing internetworks due to their manual administration.
- Static routers are not fault tolerant. The lifetime of a manually configured static route is infinite and, therefore, static routers do not sense and recover from downed routers or downed links.
- A good example of a static router is a multihomed computer running Windows 2000 (a computer with multiple network interface cards). Creating a static IP router with Windows 2000 is as simple as installing multiple network interface cards, configuring TCP/IP, and enabling IP routing.
- Top Of Page

- **Dynamic Routing**
- A router with dynamically configured routing tables is known as a dynamic router. Dynamic routing consists of routing tables that are built and maintained automatically through an ongoing communication between routers. This communication is facilitated by a routing protocol, a series of periodic or on-demand messages containing routing information that is exchanged between routers. Except for their initial configuration, dynamic routers require little ongoing maintenance, and therefore can scale to larger internetworks.
- Dynamic routing is fault tolerant. Dynamic routes learned from other routers have a finite lifetime. If a router or link goes down, the routers sense the change in the internetwork topology through the expiration of the lifetime of the learned route in the routing table. This change can then be propagated to other routers so that all the routers on the internetwork become aware of the new internetwork topology.
- The ability to scale and recover from internetwork faults makes dynamic routing the better choice for medium, large, and very large internetworks.
- A good example of a dynamic router is a computer with Windows 2000 Server and the Routing and Remote Access Service running the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) routing protocols for IP and RIP for IPX.
- ## Host Routing

- When a host using a routable protocol wants to send data to another host, it must first obtain the internetwork address of the destination. The destination internetwork address is obtained through an address resolution process whereby the sending host obtains the destination internetwork address by referencing its logical name. For example, TCP/IP hosts use Domain Name System (DNS) name resolution to resolve a DNS domain name to an IP address. Novell NetWare workstations query the bindery (a database stored on a NetWare server) or directory tree of their default server to resolve a server name to its Internetwork Packet Exchange (IPX) internetwork address.
- Once the destination internetwork address has been obtained, the source network and the destination network addresses are compared. When the source and destination hosts are on the same network, the packets are sent directly to the destination host by the source without the use of a router (see Figure 1.3). The source host sends the packet to the destination by addressing the packet to the destination's physical address. This is known as a direct delivery. In a direct delivery, the destination internetwork address and the destination physical address are for the same end system.
- Conversely, when the source and destination hosts are on different networks, the packets to the destination cannot be directly delivered by the source. Instead, the source delivers them to an intermediate router (see Figure 1.3) by addressing the packet to the router's physical address. This is known as an indirect delivery. In an indirect delivery, the destination internetwork address and the destination physical address are not for the same end system.
- During an indirect delivery, the sending host forwards the packet to a router on its network by determining the router corresponding to the first hop or by discovering the entire path from the source to the destination.

- 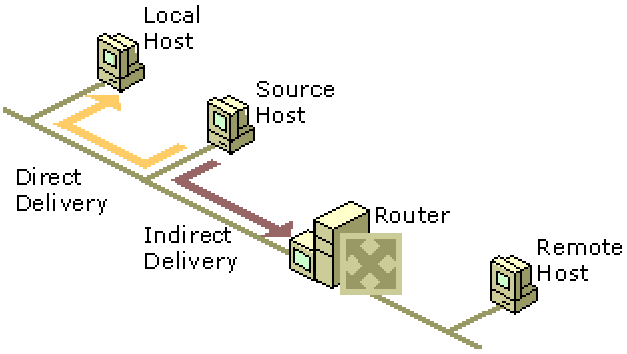

- **Figure 1.3 Host Routing Process**
- **Host Determination of the First Hop**
- IP and IPX sending hosts determine the physical address of the first hop router using one of the following processes:
- **Host routing table** A routing table on the host yields the forwarding address of the router to be used to reach the desired destination network ID. An example is the IP routing table on a Microsoft TCP/IP host. See "Routing Tables" later in this chapter for a detailed definition of a routing table.
- **Dynamic updates of host routing table** TCP/IP has a facility to dynamically update the host routing table with better routes, as packets are sent to destinations. The Internet Control Message Protocol (ICMP) Redirect message is sent by an IP router to a sending host informing it of a better route to a destination host. The better route becomes a host route in the routing table. TCP/IP for Windows 2000 supports the dynamic update of the IP routing table based on the receipt of the ICMP Redirect message.
- **Eavesdropping** TCP/IP hosts have the ability to listen to the routing protocol traffic used by routers. This is known as eavesdropping or wiretapping. Eavesdropping hosts have the same routing information as the routers. An example of eavesdropping is Silent RIP. Silent RIP is the ability of a TCP/IP host to listen to RIP for IP routing traffic exchanged by RIP routers and update its routing table. Microsoft® Windows NT® Server 3.51 and Service Pack 2 and later, Microsoft® Windows NT® Workstation 4.0 and Service Pack 4 and later support Silent RIP.
- **Default route** To simplify the configuration of hosts and routers and to reduce the overhead associated with each host having routes for all the networks in the internetwork, a sending host is configured with a single default route. The default route and its forwarding address to the default router are used when no other routes to the destination network are found. The Default Gateway for TCP/IP hosts is a default router.
- **Querying the network for the best route** For hosts without a routing table or a configured default router, the sending host can determine the physical address of the first hop router by querying the routers on the network. A query for the best route to a specified destination network address is sent as a broadcast or multicast packet. The responses from the routers are analyzed by the sending host, and the best router is chosen. An example of this querying process is the RIP GetLocalTarget message sent by an IPX host. The Routing Information Protocol (RIP) GetLocalTarget message contains a desired destination IPX network ID. IPX routers on the sending host's network that can reach the destination IPX network ID send a response to the sending host. Based on the RIP responses from the local routers, the sending host chooses the best router to forward the IPX packet.
- Top Of Page

- **Host Determination of the Entire Path**
- When using some routable protocols, the sending host does more than determine the first hop. The source host goes through a route discovery process and determines the path between the sending host and the destination. The list of networks or routers is then included in the Network layer header and is used by the routers to forward the packet along the indicated path. This process is known as source routing.
- In source routing, the routers are only acting as store and forward devices because the routing decisions have already been made by the sending host. Source routing is not typically implemented as a method of routing because the path either needs to be known or discovered. Source route discovery processes tend to be traffic intensive and slow. IP routing is normally done through routing decisions made by sending hosts and IP routers based on local routing tables. However, in network testing and debugging situations, it is sometimes desired to specify an exact route through the IP internetwork that overrides the path that would normally be taken. This is known as IP source routing.
- In IP source routing, the entire route is specified by the sending host through the IP addresses of successive IP routers between the source and destination. At each IP router, the IP datagram is addressed to the next router using the Destination IP address field of the IP header.
- IP supports two types of source routing. The first type is loose source routing, in which the IP address of the next router can be one or more routers away (multiple hops). The second type is strict source routing, in which the next router must be a neighboring router (single hop).
-

- **Note**
- Token Ring source routing is a Media Access Control (MAC)–sublayer routing scheme and does not apply to the internetwork-based source routing discussed earlier.

Routing Tables

- During the routing process, the routing decisions of hosts and routers are aided by a database of routes known as the routing table. The routing table is not exclusive to a router. Depending on the routable protocol, hosts may also have a routing table that may be used to decide the best router for the packet to be forwarded. IP hosts have a routing table. IPX hosts do not have a routing table.
- The types of possible entries in a routing table include:
- *Network Route* . A route to a specific Network ID in the internetwork.
- *Host Route* . A route to a specific internetwork address (Network ID and Host ID). Instead of making a routing decision based on just the network ID, the routing decision is based on the combination of network ID and host ID. Host routes allow intelligent routing decisions to be made for each internetwork address. Host routes are typically used to create custom routes to control or optimize specific types of internetwork traffic.
- *Default Route* . A route that is used when no other routes for the destination are found in the routing table. For example, if a router or end system cannot find a network route or host route for the destination, the default route is used. Rather than being configured with routes for all the Network IDs in the internetwork, the default route is used to simplify the configuration of end systems or routers.
- **Note**
- In many router implementations including the Windows 2000 Routing and Remote Access service, there is a routing table and a forwarding table. The routing table is used to store all the routes from all possible sources. The forwarding table is what is used by the routable protocol when forwarding the packet. For example, for a Windows 2000 router, the Routing and Remote Access service maintains the IP routing table using a component called the Route Table Manager. The IP forwarding table is contained within the TCP/IP protocol. The Route Table Manager updates the IP forwarding table based on incoming route information from multiple sources. The contents of the routing table do not necessarily match the contents of the forwarding table. For the purposes of discussion in this introductory chapter, the routing table and the forwarding table are the same.

Routing Table Structure

- As illustrated in Figure 1.5, entries in the routing table usually consist of the following fields:
- **Network ID** The Network ID field contains the identification number for a network route or an internetwork address for a host route.
- **Forwarding Address** The Forwarding Address field contains the address to which the packet is to be forwarded. The forwarding address can be a network interface card address or an internetwork address. For network IDs to which the end system or router is directly attached, the Forwarding Address field may be blank.
- **Interface** The Interface field indicates the network interface that is used when forwarding packets to the network ID. This is a port number or other type of logical identifier. For example, the interface for a 3COM EtherLink III network interface card may be referred to as ELNK3 in the routing table.
- **Metric** The Metric field indicates the cost of a route. If multiple routes exist to a given destination network ID, the metric is used to decide which route is to be taken. The route with the lowest metric is the preferred route. Some routing algorithms only store a single route to any Network ID in the routing table even when multiple routes exist. In this case, the metric is used by the router to decide which route to store in the routing table.
- Metrics can indicate different ways of expressing a route preference:
- *Hop Count* . A common metric. Indicates the number of routers (hops) in the path to the network ID.
- *Delay* . A measure of time that is required for the packet to reach the network ID. Delay is used to indicate the speed of the path—local area networks (LAN) links have a low delay, wide area network (WAN) links have a high delay—or a congested condition of a path.
- *Throughput* . The effective amount of data that can be sent along the path per second. Throughput is not necessarily a reflection of the bit rate of the link, as a very busy Ethernet link may have a lower throughput than an unutilized 64-Kbps WAN link.
- *Reliability* . A measure of the path constancy. Some types of links are more prone to link failures than others. For example, with WAN links, leased lines are more reliable than dial-up lines.
- **Lifetime** The Lifetime field indicates the lifetime that the route is considered valid. When routes are learned through the exchange of information with other routers, this is an additional field that is used. Learned routes have a finite lifetime. To keep a learned route in the routing table, the route must be refreshed through a periodic process. If a learned route's lifetime expires, it is removed from the routing table. The timing out of learned routes provides a way for routers to reconfigure themselves when the topology of an internetwork changes due to a downed link or a downed router.

- 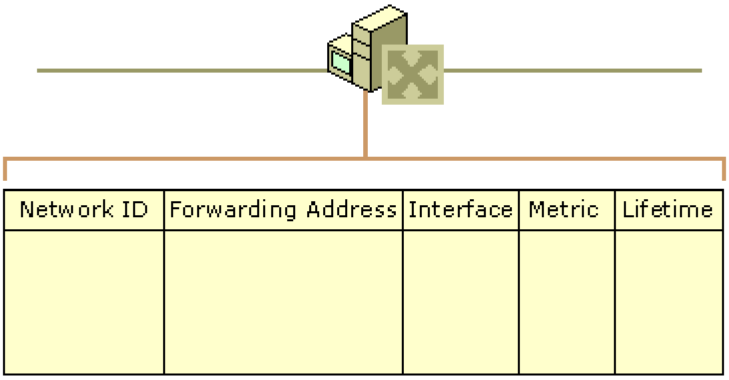

Network ID	Forwarding Address	Interface	Metric	Lifetime

- **Figure 1.5 Routing Table Structure**
- **Note**
- The Lifetime field is typically not visible in routing tables.
- This list of fields is a representative list in the routing tables. Actual fields in the routing tables for different routable protocols may vary. For information about the IP routing table, see "Introduction to TCP/IP" in the *TCP/IP Core Networking Guide* . For information about the IPX routing table, see "IPX Routing" in this book.
- Top Of Page

- **Locality of the Routing Table**
- All the routing decisions made by the end system or the router are based on information in a local routing table that physically resides in the random access memory (RAM) of the system making the routing decision. There is no single, holistic view of the internetwork that is being gathered by a server and downloaded to each end system and router so that all users have the same view of the internetwork and all traffic flows along predictable pathways.
- Each router in a path between a source and destination makes a local routing decision based on its local routing table. The path taken from the source to the destination may not be the same as the path for response packets from the destination back to the source. If the information in the local routing tables of the end systems or routers is incorrect due to misconfiguration or changing network conditions, then routing problems can result. Troubleshooting routing problems may involve the analysis of the routing tables of the end systems (source and destination) and all the routers forwarding packets between them.
- For information about the operation and troubleshooting of IP routing, see "Unicast IP Routing" in this book. For information about the operation and troubleshooting of IPX routing, see "IPX Routing" in this book.

Foundations of Routing Protocols

Dynamic routers use routing protocols to facilitate the ongoing communication and dynamic updating of routing tables. Routing protocols are used between routers and represent additional network traffic overhead on the network. This additional traffic can become an important factor in planning WAN link usage. RIP and OSPF for IP, and RIP and NLSP for IPX are all routing protocols. In some cases, such as RIP for IP (version 1) and RIP for IPX, the routing information is exchanged using MAC-level broadcasts.

An important element of a routing protocol implementation is its ability to sense and recover from internetwork faults. How quickly it can recover is determined by the type of fault, how it is sensed, and how the routing information is propagated through the internetwork.

When all the routers on the internetwork have the correct routing information in their routing tables, the internetwork has converged. When convergence is achieved, the internetwork is in a stable state and all routing occurs along optimal paths.

When a link or router fails, the internetwork must reconfigure itself to reflect the new topology. Information in routing tables must be updated. Until the internetwork reconverges, it is in an unstable state in which routing loops and black holes can occur. The time it takes for the internetwork to reconverge is known as the convergence time. The convergence time varies based on the routing protocol and the type of failure (downed link or downed router).

Routing protocols are based either on a distance vector or link state technology. The main differences between distance vector and link state routing protocols include the following:

- What routing information is exchanged.

- How the information is exchanged.
- How quickly the internetwork can recover from a downed link or a downed router.

Internetwork Routing

The following terms are essential to your understanding of routing:

End Systems . As defined by the International Standards Organization (ISO), end systems are network devices without the ability to forward packets between portions of a network. End systems are also known as hosts.

Intermediate Systems . Network devices with the ability to forward packets between portions of a network. Bridges, switches, and routers are examples of intermediate systems.

Network . A portion of the networking infrastructure (encompassing repeaters, hubs, and bridges/Layer 2 switches that is bound by a network layer intermediate system and is associated with the same network layer address.

Router . A network layer intermediate system used to connect networks together based on a common network layer protocol.

Hardware Router . A router that performs routing as a dedicated function and has specific hardware designed and optimized for routing.

Software Router . A router that is not dedicated to performing routing but performs routing as one of multiple processes running on the router computer. The Windows 2000 Server Router Service is a software router.

Internetwork . At least two networks connected using routers. Figure 1.1 illustrates an internetwork.

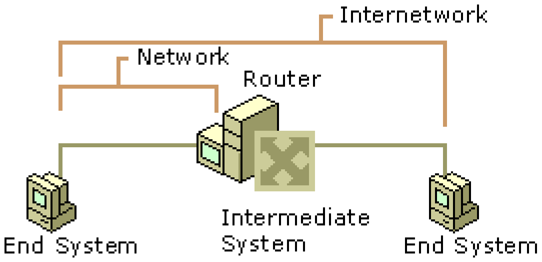

Figure 1.1 An Internetwork

Addressing in an Internetwork

The following internetwork addressing terms are also important to your understanding of routing:

Network address . Also known as a network ID. The number assigned to a single network in an internetwork. Network addresses are used by hosts and routers when routing a packet from a source to a destination in an internetwork.

Host address . Also known as a host ID or a node ID. Can either be the host's physical address (the address of the network interface card) or an administratively assigned address that uniquely identifies the host on its network.

Internetwork address . The combination of the network address and the host address; it uniquely identifies a host on an internetwork. An IP address that contains a network ID and a host ID is an internetwork address.

For detailed information about how IP implements network ID and host ID addressing, see "Introduction to TCP/IP" in the *Microsoft ® Windows ® 2000 Server Resource Kit TCP/IP Core Networking Guide* .

When a packet is sent from a source host to a destination host on an internetwork, the Network layer header of the packet contains:

- The Source Internetwork Address, which contains a source network address and source host address.
- The Destination Internetwork Address, which contains a destination network address and destination host address.
- A Hop Count, which either starts at zero and increases numerically for each router crossed to a maximum value, or starts at a maximum value and decreases numerically to zero for each router crossed. The hop count is used to prevent the packet from endlessly circulating on the internetwork.

Top Of Page

Demand-Dial Routing

Microsoft® Windows® 2000 provides extensive support for demand-dial routing, the routing of packets over point-to-point links such as analog phone lines and ISDN. Demand-dial routing allows you to connect to the Internet, to connect branch offices, or to implement router-to-router virtual private network (VPN) connections.

IPX Routing

The Microsoft® Windows® 2000 Server Router is a fully functional Internetwork Packet Exchange(IPX) router supporting Routing Information Protocol (RIP) for IPX, the primary routing protocol used in IPX internetworks; Novell NetWare Service Advertising Protocol (SAP) for IPX, a protocol for the collection and distribution of service names and addresses; and NetBIOS over IPX broadcast forwarding

MANAGING & MONITORING IP ROUTING

Managing VSANs

Table Of Contents

Monitoring SNMP Traffic

Managing VSANs

VSANs (virtual SANs) allow you to separate devices that are physically connected to the same fabric, and thus provide higher security and greater scalability in the network fabric. When you create VSANs, you are creating multiple logical SANs over a common physical infrastructure. After creating VSANs, you must establish IP static routes between the network segments if you are using the IP over Fibre Channel (IPFC) protocol to manage your Cisco MDS 9000 Family switches.

The Fabric Manager allows you to configure VSANs on multiple Cisco 9000 switches. The Device Manager allows you to configure VSANs on a single Cisco 9000 switch.

Note For information about VSANs and configuring them using the command-line interface (CLI), refer to the *Cisco 9000 Family Configuration Guide*.

You can manage Cisco MDS 9000 Family switches through Ethernet connections to the management interface (mgmt 0) of each switch or by using the IPFC protocol. To use IPFC, you connect to a switch using the Ethernet management interface and establish routes from that switch to the other switches over the Fibre Channel network. When you segment the Fibre Channel network using VSANs, you must establish static routes between the network segments.

Figure 4-1 shows a physical Fibre Channel network with two VSANs. VSAN 2 is connected by dashed lines and VSAN 7 is connected by solid lines.

Figure 4-1 Configuring VSANs

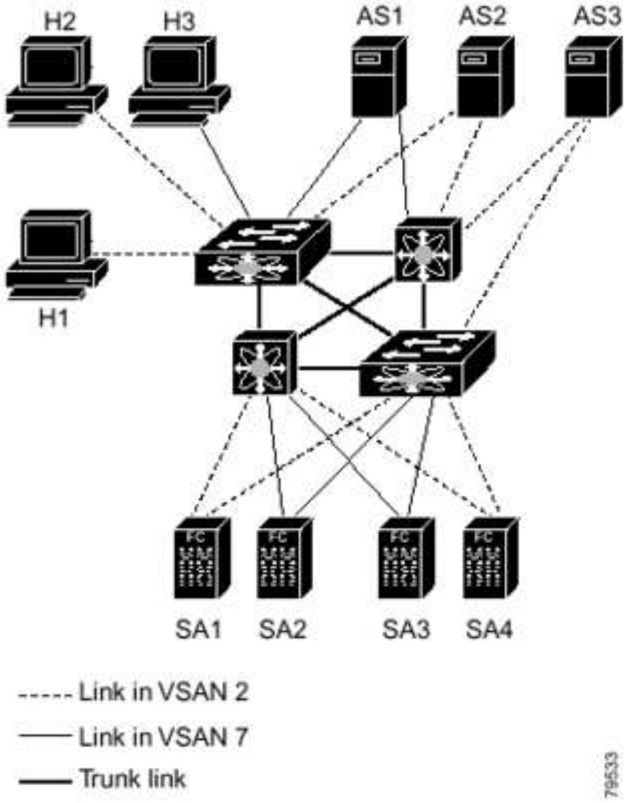

VSAN 2 includes the H1 and H2 hosts, the AS2 and AS3 application servers, and the SA1 and SA4 storage arrays. VSAN 7 connects H3, AS1, SA2, and SA3. The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic.

VSAN 1 is the default VSAN for Cisco MDS 9000 Family switches. All ports are assigned by default to VSAN 1. VSAN 4094 is called the isolated VSAN. When a VSAN is deleted, any ports in that VSAN are moved to VSAN 4094.

Note We recommend that you delete or move all the ports in a VSAN before deleting the VSAN.

VSANs are enabled through trunking, which enables interconnect ports to transmit and receive frames in more than one VSAN over a single physical link, using the Extended Inter-Switch Link (EISL) protocol. The trunking protocol is enabled by default, and if disabled on a switch, no ports on that switch or directly connected to the switch will support the use of VSANs.

By default, the trunk mode is enabled on all Fibre Channel interfaces, but can be disabled on a port-by-port basis. When connected to a third-party switch, the trunk mode configuration has no effect—the ISL is always in a trunking disabled state.

Each Fibre Channel interface has an associated trunk-allowed VSAN list. This list determines the VSANs that are supported on each interface. By default, the entire range of VSANs from 1 through 4093 are allowed on any interface. You can restrict an interface to the use of a specific set of VSANs, which prevents traffic from any other VSAN being transmitted on the interface.

Procedures for managing VSANs include:

- Adding and Configuring VSANs

Adding and Configuring VSANs

To add and configure VSANs, perform the following steps.

Step 1 From the Fabric Manager, choose **FC > VSAN** from the menu tree, OR
From Device Manager, choose the **VSAN** option from the FC menu or click the **VSAN** icon on the toolbar.

The Fabric Manager's Information pane displays VSAN attributes for multiple switches. The VSAN dialog box in the Device Manager displays VSAN general attributes for a single switch.

Step 2 From Fabric Manager, click **Create Row** on the Information pane toolbar, OR
From Device Manager, click **Create** on the VSAN dialog box.

You see the Create dialog box.

Step 3 Complete the fields on this dialog box and click **OK** to add the VSAN.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Controlling In-Band Management Connectivity

The Fabric Manager allows you to configure and monitor IP traffic on multiple Cisco MDS 9000 Family switches. The Device Manager allows you to configure and monitor IP traffic on a single Cisco 9000 switch.

Cisco MDS 9000 Family switches support both out-of-band and in-band management schemes. An Ethernet connection provides out-of-band management using Telnet, SSH or SNMP access. In-band IP management is also available using IP over Fibre Channel (IPFC). IPFC encapsulates IP packets into Fibre Channel frames so that management information can cross the Fibre Channel network without requiring a dedicated Ethernet connection to each switch. IP addresses are resolved to the Fibre Channel address through the Address Resolution Protocol (ARP).

Procedures for managing and viewing connectivity information include:

- Configuring IP Routing for Management Traffic

- Managing IPFC Connectivity with Multiple VSANs
- Viewing IP Address Information
- Enabling or Disabling IP Forwarding
- Viewing TCP Information and Statistics
- Viewing UDP Information and Statistics
- Viewing IP Statistics
- Viewing ICMP Statistics
- Monitoring SNMP Traffic

Configuring IP Routing for Management Traffic

When using in-band network management over Fibre Channel links, you must ensure that a path exists from the seed switch, connected to the Cisco Fabric Manager over its Ethernet interface (mgmt0), and the other switches in the network fabric. See Figure 4-2.

Figure 4-2 IP Routing Between VSANs

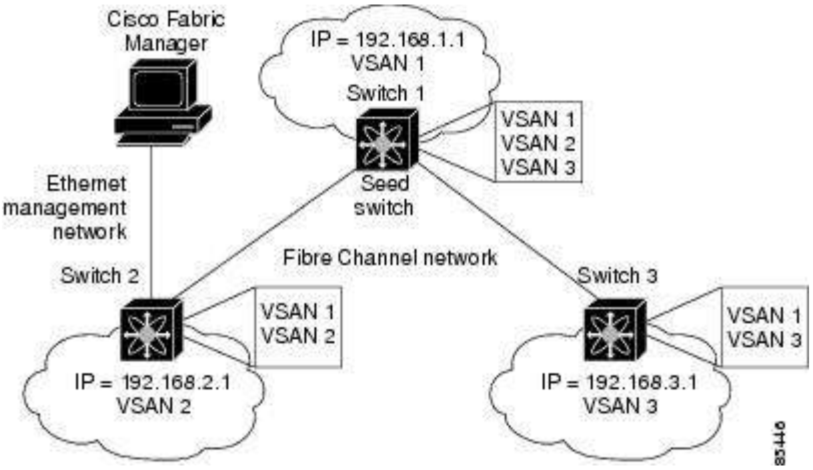

To do this, make sure that the seed switch has a path to each VSAN. Each of the other switches can then be configured to use the seed switch as their default gateway. For example, in Figure 4-2, switch 1 is connected to VSAN 2 and VSAN 3, while switch 2 and switch 3 are configured to use switch 1 as their default gateway.

You can also configure static routes on a point-to-point basis from one switch to another. In this example, you would configure a static route on both switch 2 and switch 3 to switch 1.

Configuring an IP Route

To configure an IP route or identify the default gateway, perform the following steps.

Step 1 From the Device Manager, choose **Routes** from the **IP** menu.

You see the IP Routes window.

Step 2 To create a new IP route or identify the default gateway on a switch, click the **Create** button.

You see the Create IP Routes window.

Step 3 Complete the fields on this window and click **OK** to add an IP route.

Step 4 To configure a static route, enter the destination network ID and subnet mask in the Dest and Mask fields. To configure a default gateway, enter the IP address of the seed switch in the Gateway field.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Managing IPFC Connectivity with Multiple VSANs

To configure IPFC from the Device Manager, choose **VSAN** from the FC menu and click the **General** tab.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Viewing IP Address Information

To view IP addresses of the switches in the current fabric from the Fabric Manager, choose **Switches** from the menu tree.

The Information pane displays IP address information for multiple switches.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Enabling or Disabling IP Forwarding

To view or change the IP forwarding configuration of the switches in the current fabric, perform the following steps.

Step 1 Choose **IP > Forwarding** from the Fabric Manager menu tree.

Step 2 To enable IP forwarding for a specific switch, click the **RoutingEnabled** check box.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Viewing TCP Information and Statistics

To view TCP information from the Device Manager, choose **Mgmt TCP/UDP** from the IP menu.

To monitor TCP statistics from the Fabric Manager, choose IP > Mgmt Statistics from the menu tree and click the TCP tab. To monitor TCP statistics from the Device Manager, choose Statistics from the IP menu and view the TCP tab.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Viewing UDP Information and Statistics

To view User Datagram Protocol (UDP) information, from the Device Manager, choose **Mgmt TCP/UDP** from the IP menu and click the **UDP** tab.

To monitor UDP traffic from the Fabric Manager, choose **IP** > Mgmt Statistics from the menu tree and click the UDP tab. From Device Manager, choose **Statistics** from the IP menu and click the UDP tab.

The Fabric Manager Information pane displays TCP traffic information for multiple switches. The Device Manager dialog box displays information for a single switch.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Viewing IP Statistics

To monitor IP statistics from the Fabric Manager, choose **IP** >Mgmt **Statistics** from the menu tree and click the **IP** tab. From Device Manager, select **Statistics** from the IP menu and click the IP tab.

The Fabric Manager Information pane displays IP statistics for multiple switches. The Device Manager dialog box displays information for a single switch.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Viewing ICMP Statistics

To monitor statistics for ICMP packets received, select **IP** > Mgmt **Statistics** from the menu tree and click the **ICMP In** tab. To monitor statistics for ICMP packets transmitted from the Fabric Manager, select **IP** > Mgmt **Statistics** from the menu tree and click the **ICMP Out** tab.

To monitor ICMP statistics from Device Manager, select Statistics from the IP menu and click the **ICMP** tab.

The Fabric Manager Information pane displays information for multiple switches. The Device Manager dialog box displays information for a single switch.

In the Device Manager, a prefix (In or Out) identifies whether the packets are received or transmitted. In the Fabric Manager, separate tabs on the Information pane are provided for incoming and outbound ICMP traffic and this prefix is omitted.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems.

Monitoring SNMP Traffic

To monitor SNMP traffic statistics from the Fabric Manager, select **IP** >Mgmt **Statistics** from the menu tree and click on the **SNMP** tab. To monitor SNMP traffic from Device Manager, select **Statistics** from the IP menu and click the **SNMP** tab.

The Fabric Manager Information pane displays information for multiple switches. The Device Manager dialog box displays information for a single switch.

Note You can access the field descriptions for the windows or dialog boxes in this procedure in the Reference section of the Fabric Manager or Device Manager help systems

Network address translation

From Wikipedia, the free encyclopedia
Jump to: navigation, search
"NAT" redirects here. For other uses, see Nat (disambiguation).

In computer networking, **network address translation** (**NAT**) is the process of modifying network address information in datagram (IP) packet headers while in transit across a traffic routing device for the purpose of remapping one IP address space into another.

Most often today, NAT is used in conjunction with *network masquerading* (or *IP masquerading*) which is a technique that hides an entire IP address space, usually consisting of private network IP addresses (RFC 1918), behind a single IP address in another, often public address space. This mechanism is implemented in a routing device that uses stateful translation tables to map the "hidden" addresses into a single IP address and readdresses the outgoing Internet Protocol (IP) packets on exit so that they appear to originate from the router. In the reverse communications path, responses are mapped back to the originating IP address using the rules ("state") stored in the translation tables. The translation table rules established in this fashion are flushed after a short period unless new traffic refreshes their state.

As described, the method enables communication through the router only when the conversation originates in the masqueraded network, since this establishes the translation tables. For example, a web browser in the masqueraded network can browse a website outside, but a web browser outside could not browse a web site in the masqueraded network. However, most NAT devices today allow the network administrator to configure translation table entries for permanent use. This feature is often referred to as "static NAT" or port forwarding and allows traffic originating in the "outside" network to reach designated hosts in the masqueraded network.

Because of the popularity of this technique (see below), the term *NAT* has become virtually synonymous with the method of IP masquerading.

Network address translation has serious consequences, both drawbacks and benefits, on the quality of Internet connectivity and requires careful attention to the details of its implementation. As a result, many methods have been devised to alleviate the issues encountered. See the article on *NAT traversal*.

Contents

[hide]

[edit] Overview

In the mid-1990s (1994) NAT became a popular tool for alleviating the problem of IPv4 address exhaustion. It has become a standard, indispensable feature in routers for home and small-office Internet connections.

Most systems using NAT do so in order to enable multiple hosts on a private network to access the Internet using a single public IP address (see gateway). However, NAT breaks the originally envisioned model of IP end-to-end connectivity across the Internet, introduces complications in communication between hosts, and affects performance.

NAT obscures an internal network's structure: all traffic appears to outside parties as if it originated from the gateway machine.

Network address translation involves over-writing the source or destination IP address and usually also the TCP/UDP port numbers of IP packets as they pass through the router. Checksums (both IP and TCP/UDP) must also be rewritten as a result of these changes.

In a typical configuration, a local network uses one of the designated "private" IP address subnets (RFC 1918). Private Network Addresses are 192.168.x.x, 172.16.x.x through 172.31.x.x, and 10.x.x.x (or using CIDR notation, 192.168/16, 172.16/12, and 10/8), and a router on that network has a private address (such as 192.168.0.1) in that address space. The router is also connected to the Internet with a single "public" address (known as "overloaded" NAT) or multiple "public" addresses assigned by an ISP. As traffic passes from the local network to the Internet, the source address in each packet is translated on the fly from the private addresses to the public address(es). The router tracks basic data about each active connection (particularly the destination address and port). When a reply returns to the router, it uses the connection tracking data it stored during the outbound phase to determine where on the internal network to forward the reply; the TCP or UDP client port numbers are used to demultiplex the packets in the case of overloaded NAT, or IP address and port number when multiple public addresses are available, on packet return. To a host on the Internet, the router itself appears to be the source/destination for this traffic.

[edit] Basic NAT and PAT

There are two levels of network address translation.

- **Basic NAT**. This involves IP address translation only, not port mapping.
- **PAT (Port Address Translation)**. Also called simply "NAT" or "Network Address Port Translation, NAPT". This involves the translation of both IP addresses and port numbers.

All Internet packets have a source IP address and a destination IP address. Both or either of the source and destination addresses may be translated.

Some Internet packets do not have port numbers: for example, ICMP packets. However, the vast bulk of Internet traffic is TCP and UDP packets, which do have port numbers. Packets which do have port numbers have both a source port number and a destination port number. Both or either of the source and destination ports may be translated.

NAT which involves translation of the source IP address and/or source port is called **source NAT** or **SNAT**. This re-writes the IP address and/or port number of the computer which originated the packet.

NAT which involves translation of the destination IP address and/or destination port number is called **destination NAT** or **DNAT**. This re-writes the IP address and/or port number corresponding to the destination computer.

SNAT and DNAT may be applied simultaneously to Internet packets.

[edit] Types of NAT

Network address translation is implemented in a variety of schemes of translating addresses and port numbers, each affecting application communication protocols differently. In some application protocols that use IP address information, the application running on a node in the masqueraded network needs to determine the external address of the NAT, i.e., the address that its communication peers detect, and, furthermore, often needs to examine and categorize the type of mapping in use. For this purpose, the Simple traversal of UDP over NATs (STUN) protocol was developed (RFC 3489, March 2003). It classified NAT implementation as *full cone NAT, (address) restricted cone NAT*, *port restricted cone NAT* or *symmetric NAT* and proposed a methodology for testing a device accordingly. However, these procedures have since been deprecated from standards status, as the methods have proven faulty and inadequate to correctly assess many devices. New methods have been standardized in RFC 5389 (October 2008) and the STUN acronym now represents the new title of the specification: *Session Traversal Utilities for NAT*.

Full-cone NAT, also known as *one-to-one NAT*

- Once an internal address (iAddr:iPort) is mapped to an external address (eAddr:ePort), any packets from iAddr:iPort will be sent through eAddr:ePort.
- *Any external host* can send packets to iAddr:iPort by sending packets to eAddr:ePort.

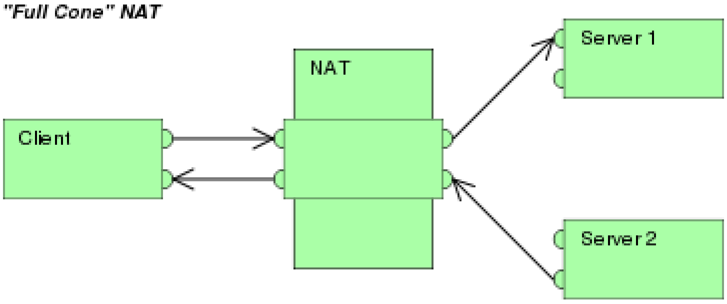

(Address) restricted cone NAT

- Once an internal address (iAddr:iPort) is mapped to an external address (eAddr:ePort), any packets from iAddr:iPort will be sent through eAddr:ePort.
- An external host (*hAddr:any*) can send packets to iAddr:iPort by sending packets to eAddr:ePort only if iAddr:iPort has previously sent a packet to hAddr:*any*. "Any" means the port number doesn't matter.

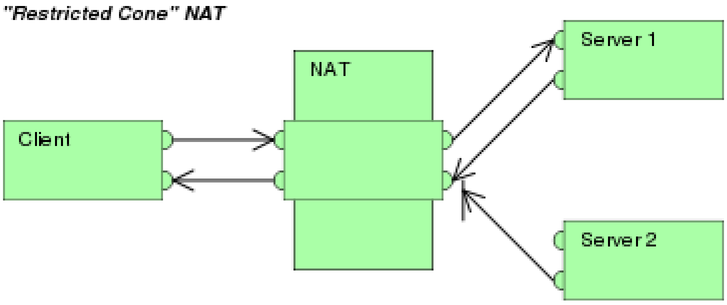

Port-restricted cone NAT

Like an address restricted cone NAT, but the restriction includes port numbers.

- Once an internal address (iAddr:iPort) is mapped to an external address (eAddr:ePort), any packets from iAddr:iPort will be sent through eAddr:ePort.
- An external host (*hAddr:hPort*) can send packets to iAddr:iPort by sending packets to eAddr:ePort only if iAddr:iPort has previously sent a packet to hAddr:hPort.

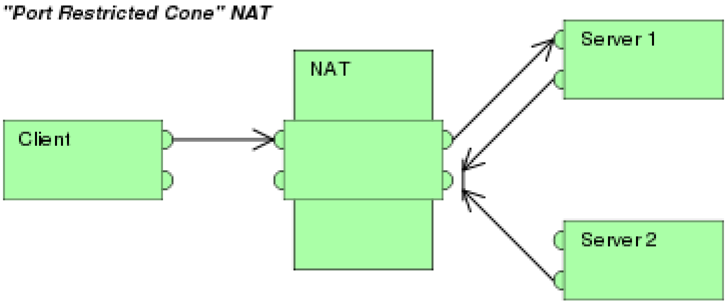

Symmetric NAT

- Requests from internal IP address and port combinations to different external IP address and port pairs are mapped to the external NAT address on a unique port. This also applies to all requests from the same host to different destinations.
- Only an external host that receives a packet from an internal host can send a packet back.

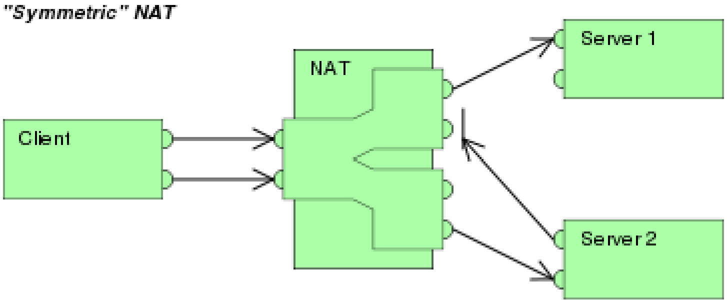

This terminology has been the source of much confusion, as it has proven inadequate at describing real-life NAT behavior.[1] Many NAT implementations combine these types, and it is therefore better to refer to specific individual NAT behaviors instead of using the

Cone/Symmetric terminology. Especially, most NAT translators combine *symmetric NAT* for outgoing connections with *static port mapping*, where incoming packets to the external address and port are redirected to a specific internal address and port. Some products can redirect packets to several internal hosts, e.g. to divide the load between a few servers. However, this introduces problems with more sophisticated communications that have many interconnected packets, and thus is rarely used.

Many NAT implementations follow the *port preservation* design. For most communications, they use the same values as internal and external port numbers. However, if two internal hosts attempt to communicate with the same external host using the same port number, the external port number used by the second host will be chosen at random. Such NAT will be sometimes perceived as *(address) restricted cone NAT* and other times as *symmetric NAT*.

[edit] NAT and TCP/UDP

"Pure NAT", operating on IP alone, may or may not correctly parse protocols that are totally concerned with IP information, such as ICMP, depending on whether the payload is interpreted by a host on the "inside" or "outside" of translation. As soon as the protocol stack is climbed, even with such basic protocols as TCP and UDP, the protocols will break unless NAT takes action beyond the network layer.

IP has a checksum in each packet header, which provides error detection only for the header. IP datagrams may become fragmented and it is necessary for a NAT to reassemble these fragments to allow correct recalculation of higher-level checksums and correct tracking of which packets belong to which connection.

The major transport layer protocols, TCP and UDP, have a checksum that covers all the data they carry, as well as the TCP/UDP header, plus a "pseudo-header" that contains the source and destination IP addresses of the packet carrying the TCP/UDP header. For an originating NAT to successfully pass TCP or UDP, it must recompute the TCP/UDP header checksum based on the translated IP addresses, not the original ones, and put that checksum into the TCP/UDP header of the first packet of the fragmented set of packets. The receiving NAT must recompute the IP checksum on every packet it passes to the destination host, and also recognize and recompute the TCP/UDP header using the retranslated addresses and pseudo-header. This is not a completely solved problem. One solution is for the receiving NAT to reassemble the entire segment and then recompute a checksum calculated across all packets.

Originating host may perform Maximum transmission unit (MTU) path discovery (RFC 1191) to determine the packet size that can be transmitted without fragmentation, and then set the "don't fragment" bit in the appropriate packet header field.

[edit] Destination network address translation (DNAT)

DNAT is a technique for transparently changing the destination IP address of an en-route packet and performing the inverse function for any replies. Any router situated between two endpoints can perform this transformation of the packet.

DNAT is commonly used to publish a service located in a private network on a publicly accessible IP address. This use of DNAT is also called port forwarding.

[edit] SNAT

The meaning of the term SNAT varies by vendor. Many vendors have proprietary definitions for SNAT. A common definition is Source NAT, the counterpart of Destination NAT (DNAT). Microsoft uses the acronym for Secure NAT, in regard to the ISA Server extension discussed below. For Cisco Systems, SNAT means Stateful NAT.

The Internet Engineering Task Force (IETF) defines SNAT as Softwires Network Address Translation. This type of NAT is named after the Softwires working group that is charged with the standardization of discovery, control and encapsulation methods for connecting IPv4 networks across IPv6 networks and IPv6 networks across IPv4 networks.

[edit] Dynamic network address translation

Dynamic NAT, just like static NAT, is not common in smaller networks but is found within larger corporations with complex networks. The way dynamic NAT differs from static NAT is that where static NAT provides a one-to-one internal to public static IP address mapping, dynamic NAT doesn't make the mapping to the public IP address static and usually uses a group of available public IP addresses.

[edit] Applications affected by NAT

Some Application Layer protocols (such as FTP and SIP) send explicit network addresses within their application data. FTP in active mode, for example, uses separate connections for control traffic (commands) and for data traffic (file contents). When requesting a file transfer, the host making the request identifies the corresponding data connection by its network layer and transport layer addresses. If the host making the request lies behind a simple NAT firewall, the translation of the IP address and/or TCP port number makes the information received by the server invalid. The Session Initiation Protocol (SIP) controls Voice over IP (VoIP) and suffers the same problem. SIP may use multiple ports to set up a connection and transmit voice stream via RTP. IP addresses and port numbers are encoded in the payload data and must be known prior to the traversal of NATs. Without special techniques, such as STUN, NAT behavior is unpredictable and communications may fail.

Application Layer Gateway (ALG) software or hardware may correct these problems. An ALG software module running on a NAT firewall device updates any payload data made invalid by address translation. ALGs obviously need to understand the higher-layer protocol that they need to fix, and so each protocol with this problem requires a separate ALG.

Another possible solution to this problem is to use NAT traversal techniques using protocols such as STUN or ICE, or proprietary approaches in a session border controller. NAT traversal is possible in both TCP- and UDP-based applications, but the UDP-based technique is simpler, more widely understood, and more compatible with legacy NATs. In either case, the high level protocol must be designed with NAT traversal in mind, and it does not work reliably across symmetric NATs or other poorly-behaved legacy NATs.

Other possibilities are UPnP (Universal Plug and Play) or Bonjour (NAT-PMP), but these require the cooperation of the NAT device.

Most traditional client-server protocols (FTP being the main exception), however, do not send layer 3 contact information and therefore do not require any special treatment by NATs. In fact, avoiding NAT complications is practically a requirement when designing new higher-layer protocols today.

NATs can also cause problems where IPsec encryption is applied and in cases where multiple devices such as SIP phones are located behind a NAT. Phones which encrypt their signaling with IPsec encapsulate the port information within the IPsec packet meaning that NA(P)T devices cannot access and translate the port. In these cases the NA(P)T devices revert to simple NAT operation. This means that all traffic returning to the NAT will be mapped onto one client causing the service to fail. There are a couple of solutions to this problem: one is to use TLS, which operates at level 4 in the OSI Reference Model and therefore does not mask the port number; another is to Encapsulate the IPsec within UDP - the latter being the solution chosen by TISPAN to achieve secure NAT traversal.

The DNS protocol vulnerability announced by Dan Kaminsky on 2008 July 8 is indirectly affected by NAT port mapping. To avoid DNS server cache poisoning, it is highly desirable to not translate UDP source port numbers of outgoing DNS requests from a DNS server which is behind a firewall which implements NAT. The recommended work-around for the DNS vulnerability is to make all caching DNS servers use randomized UDP source ports. If the NAT function de-randomizes the UDP source ports, the DNS server will be made vulnerable.

[edit] Drawbacks

Hosts behind NAT-enabled routers do not have end-to-end connectivity and cannot participate in some Internet protocols. Services that require the initiation of TCP connections from the outside network, or stateless protocols such as those using UDP, can be disrupted. Unless the NAT router makes a specific effort to support such protocols, incoming packets cannot reach their destination. Some protocols can accommodate one instance of NAT between participating hosts ("passive mode" FTP, for example), sometimes with the assistance of an application-level gateway (see below), but fail when both systems are separated from the Internet by NAT. Use of NAT also complicates tunneling protocols such as IPsec because NAT modifies values in the headers which interfere with the integrity checks done by IPsec and other tunneling protocols.

End-to-end connectivity has been a core principle of the Internet, supported for example by the Internet Architecture Board. Current Internet architectural documents observe that NAT is a violation of the End-to-End Principle, but that NAT does have a valid role in careful design.[2] There is considerably more concern with the use of IPv6 NAT, and many IPv6 architects believe IPv6 was intended to remove the need for NAT.[3]

Because of the short-lived nature of the stateful translation tables in NAT routers, devices on the internal network lose IP connectivity typically within a very short period of time unless they implement NAT keep-alive mechanisms by frequently accessing outside hosts. This dramatically shortens the power reserves on battery-operated hand-held devices and has thwarted more widespread deployment of such IP-native Internet-enabled devices.

Some Internet service providers (ISPs), especially in Russia, Asia and other "developing" regions provide their customers only with "local" IP addresses, due to a limited number of external IP addresses allocated to those entities.[*citation needed*] Thus, these customers must access services external to the ISP's network through NAT. As a result, the customers cannot achieve true end-to-end connectivity, in violation of the core principles of the Internet as laid out by the Internet Architecture Board.

[edit] Benefits

The primary benefit of IP-masquerading NAT is that it has been a practical solution to the impending exhaustion of IPv4 address space. Even large networks can be connected to the Internet with as little as a single IP address. The more common arrangement is having machines that require end-to-end connectivity supplied with a routable IP address, while having machines that do not provide services to outside users behind NAT with only a few IP addresses used to enable Internet access.

Some[4] have also called this exact benefit a major drawback, since it delays the need for the implementation of IPv6, :

"[…] it is possible that its [NAT's] widespread use will significantly delay the need to deploy IPv6. […] It is probably safe to say that networks would be better off without NAT […]"

[edit] Examples of NAT software

- iptables: the Linux packet filter and NAT (interface for NetFilter)
- IPFilter: Solaris, NetBSD, FreeBSD, xMach.
- PF (firewall): The OpenBSD Packet Filter.
- Netfilter Linux packet filter framework
- Internet Connection Sharing (ICS): Windows NAT+DHCP since W98SE
- WinGate: like ICS plus lots of control

[edit] See also

- AYIYA (IPv6 over IPv4 UDP thus working IPv6 tunneling over most NATs)
- Carrier Grade NAT
- Firewall
- Gateway
- Internet Gateway Device (IGD) Protocol: UPnP NAT-traversal method
- Middlebox
- NAT-PT
- Port forwarding

- Private IP address
- Proxy server
- Routing
- Subnet
- Teredo tunneling: NAT traversal using IPv6

[edit] References

1. ^ François Audet; and Cullen Jennings (January 2007) (text). *RFC 4787 Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*. IETF. http://www.ietf.org/rfc/rfc4787.txt. Retrieved 2007-08-29.
2. ^ R. Bush; and D. Meyer; RFC 3439, *Some Internet Architectural Guidelines and Philosophy*, December 2002
3. ^ G. Van de Velde *et al.*; RFC 4864, *Local Network Protection for IPv6*, May 2007
4. ^ Larry L. Peterson; and Bruce S. Davie; *Computer Networks: A Systems Approach*, Morgan Kaufmann, 2003, pp. 328-330, ISBN 1-55860-832-X

[edit] External links

- NAT-Traversal Test and results
- Characterization of different TCP NATs – Paper discussing the different types of NAT
- Anatomy: A Look Inside Network Address Translators – Volume 7, Issue 3, September 2004
- Jeff Tyson, HowStuffWorks: *How Network Address Translation Works*
- NAT traversal techniques

Works

by Jeff Tyson

- Print
- Cite
- Feedback

  Share

- Recommend

Cite This!
Close
Please copy/paste the following text to properly cite this HowStuffWorks article:

Tyson, Jeff. "How Network Address Translation Works" 02 February 2001. HowStuffWorks.com. <http://www.howstuffworks.com/nat.htm> 26 October 2010.

How Network Address Translation Works

Inside this Article

RAM Videos

- More Tech Videos »

Computer Networking Image Gallery

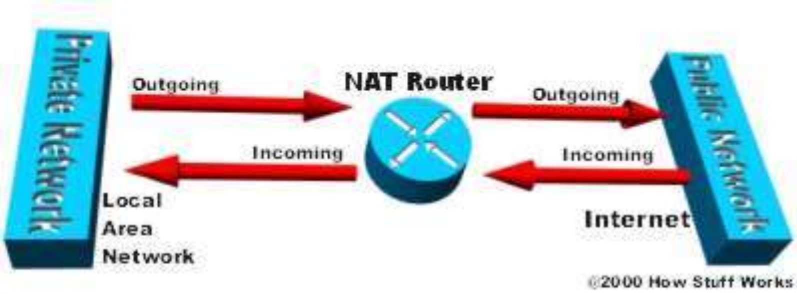

Network Address Translation helps improve security by reusing IP addresses. The NAT router translates traffic coming into and leaving the private network. See more pictures of computer networking.

If you are reading this article, you are most likely connected to the Internet and viewing it at the HowStuffWorks Web site. There's a very good chance that you are using **Network Address Translation** (NAT) right now.

The Internet has grown larger than anyone ever imagined it could be. Although the exact size is unknown, the current estimate is that there are about 100 million hosts and more than 350 million users actively on the Internet. That is more than the entire population of the United

States! In fact, the rate of growth has been such that the Internet is effectively doubling in size each year.

So what does the size of the Internet have to do with NAT? Everything! For a computer to communicate with other computers and Web servers on the Internet, it must have an **IP address**. An IP address (IP stands for Internet Protocol) is a unique 32-bit number that identifies the location of your computer on a network. Basically, it works like your street address -- as a way to find out exactly where you are and deliver information to you.

When IP addressing first came out, everyone thought that there were plenty of addresses to cover any need. Theoretically, you could have 4,294,967,296 unique addresses (2^{32}). The actual number of available addresses is smaller (somewhere between 3.2 and 3.3 billion) because of the way that the addresses are separated into classes, and because some addresses are set aside for multicasting, testing or other special uses.

With the explosion of the Internet and the increase in home networks and business networks, the number of available IP addresses is simply not enough. The obvious solution is to redesign the address format to allow for more possible addresses. This is being developed (called **IPv6**), but will take several years to implement because it requires modification of the entire infrastructure of the Internet.

This is where NAT (RFC 1631) comes to the rescue. Network Address Translation allows a single device, such as a router, to act as an agent between the Internet (or "public network") and a local (or "private") network. This means that only a single, unique IP address is required to represent an entire group of computers.

But the shortage of IP addresses is only one reason to use NAT. In this edition of **HowStuffWorks**, you will learn more about how NAT can benefit you. But first, let's take a closer look at NAT and exactly what it can do...

Network Address Translation allows a single device, such as a router, to act as agent between the Internet (or ""public network"") and a local (or ""private"") network.

Technology Information

Technology Q&A

(All IP Addressing Services Technology Q&A)

Network Address Translation (NAT) FAQ

Technology White Paper

(All IP Addressing Services Technology White Paper)

Cisco IOS NAT - Integration with MPLS VPN

Design

Design TechNotes

(All IP Addressing Services Design TechNotes)

How NAT Handles ICMP Fragments

How NAT Works

NAT: Local and Global Definitions

Network Address Translation Catalyst Switch Support Matrix

Configure

Configuration Examples and TechNotes

(All IP Addressing Services Configuration Examples and TechNotes)

Auth-proxy Authentication Inbound (Cisco IOS Firewall - Routers/Switches and NAT) Configuration Example →

Auth-proxy Authentication Inbound with IPsec and VPN Client Configuration with NAT and Cisco IOS Firewall →

Auth-proxy Authentication Outbound (Cisco IOS Firewall and NAT) Configuration →

Authentication Proxy Authentication Outbound - No Cisco IOS Firewall or NAT Configuration →

Configuring a Cisco 827 for PPPoE with VPN IPSec NAT Overloading →

Configuring a Cisco 827 Router With IRB and NAT Using RFC1483 Bridging →

Configuring a Cisco 1700/2600/3600 ADSL WIC as a PPPoE Client With NAT →

Configuring a Cisco 1700/2600/3600 ADSL WIC With NAT, a DHCP Server, and Easy IP Using PPPoA (aal5mux ppp) →

Configuring a Cisco SOHO77 Router as a PPPoE Client With NAT →

Configuring a Cisco SOHO77 Router With NAT, a DHCP Server, and Easy IP Using PPPoA (aal5mux ppp) →

Configuring a Router IPsec Tunnel Private-to-Private Network with NAT and a Static →

Configuring and Troubleshooting a Cisco 1700 Router Using an Ethernet WIC as a PPPoE Client with NAT →

Configuring Data-Link Switching and Network Address Translation →

Configuring Dynamic Multipoint VPN Using GRE Over IPSec With EIGRP, NAT, and CBAC →

Configuring Dynamic Multipoint VPN Using GRE Over IPsec With OSPF, NAT, and Cisco IOS Firewall →

Configuring IDSL on a Cisco 802/804 Router With mux-ppp, a 6160 DSLAM, Terminating on a 6400 UAC With NAT →

Configuring IPSec Router-to-Router with NAT Overload and Cisco Secure VPN Client →

Configuring IPSec Router-to-Router, Pre-shared, NAT Overload Between a Private and a Public Network →

Configuring IPSec Router-to-Router, Pre-shared, NAT Overload Between Private Networks →

Configuring Multiple VPN Clients to a Cisco VPN 3000 Concentrator Using NAT-Traversal →

Configuring NAT Transparent Mode for IPSec on the VPN 3000 Concentrator →

Configuring Network Address Translation: Getting Started

Configuring PIX to PIX Dynamic-to-Static IPSec with NAT and Cisco VPN Client →

Configuring Router to VPN Client, Mode-Config, Wild-Card Pre-Shared Key with NAT →

Configuring Router-to-Router IPsec (Pre-shared Keys) on GRE Tunnel with IOS Firewall and NAT →

Configuring Static and Dynamic NAT Simultaneously

How Does Multicast NAT Work on Cisco Routers?

How to Change the Dynamic NAT Configuration

IOS NAT Load-Balancing for Two ISP Connections

IOS NAT Load-Balancing with Optimized Edge Routing for Two Internet Connections

IPSec/GRE with NAT on IOS Router Configuration Example →

NAT in Catalyst 6500/6000 Switches Configuration Example →

NAT Order of Operation

NAT Pools and Subnet Zero

NAT Support for Multiple Pools Using Route Maps

Network Address Translation on a Stick

PIX 6.x : IPsec Tunnel Pass Through a PIX Firewall With use of Access List and with NAT Configuration Example →

Sample Configuration Using the ip nat outside source list Command

Sample Configuration Using the ip nat outside source static Command

Two-Interface Router with NAT Cisco IOS Firewall Configuration →

Using NAT and PAT Statements on the Cisco Secure PIX Firewall →

Using NAT in Overlapping Networks

Using Non-Standard FTP Port Numbers with NAT

Troubleshoot and Alerts

Troubleshooting TechNotes

(All IP Addressing Services Troubleshooting TechNotes)

Avoiding Routing Loops When Using Dynamic NAT

Configuring Data-Link Switching and Network Address Translation →

Configuring NAT Transparent Mode for IPSec on the VPN 3000 Concentrator →

Configuring Network Address Translation: Getting Started

Configuring Static and Dynamic NAT Simultaneously

How Does Multicast NAT Work on Cisco Routers?

How NAT Handles ICMP Fragments

How NAT Works

How to Change the Dynamic NAT Configuration

NAT Order of Operation

NAT Support for Multiple Pools Using Route Maps

NAT: Local and Global Definitions

Network Address Translation Catalyst Switch Support Matrix

Network Address Translation on a Stick

Using NAT and PAT Statements on the Cisco Secure PIX Firewall →

Using Non-Standard FTP Port Numbers with NAT

Verifying NAT Operation and Basic NAT Troubleshooting

NAT: Local and Global Definitions

Contents

Introduction

This document defines and clarifies the Network Address Translation (NAT) terms of inside local, inside global, outside local, and outside global.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

Refer to the Cisco Technical Tips Conventions for more information on document conventions.

Term Definitions

Cisco defines these terms as:

- **Inside local address**—The IP address assigned to a host on the inside network. This is the address configured as a parameter of the computer OS or received via dynamic address allocation protocols such as DHCP. The address is likely not a legitimate IP address assigned by the Network Information Center (NIC) or service provider.
- **Inside global address**—A legitimate IP address assigned by the NIC or service provider that represents one or more inside local IP addresses to the outside world.
- **Outside local address**—The IP address of an outside host as it appears to the inside network. Not necessarily a legitimate address, it is allocated from an address space routable on the inside.
- **Outside global address**—The IP address assigned to a host on the outside network by the host owner. The address is allocated from a globally routable address or network space.

These definitions still leave a lot to be interpreted. For this example, this document redefines these terms by first defining local address and global address. Keep in mind that the terms inside and outside are NAT definitions. Interfaces on a NAT router are defined as inside or outside with the NAT configuration commands, **ip nat inside** and **ip nat outside**. Networks to which these interfaces connect can then be thought of as inside networks or outside networks, respectively.

- **Local address**—A local address is any address that appears on the inside portion of the network.
- **Global address**—A global address is any address that appears on the outside portion of the network.

Packets sourced on the inside portion of the network have an inside local address as the source address and an outside local address as the destination address of the packet, while the packet resides on the inside portion of the network. When that same packet gets switched to the outside network, the source of the packet is now known as the inside global address and the destination of the packet is known as the outside global address.

Conversely, when a packet is sourced on the outside portion of the network, while it is on the outside network, its source address is known as the outside global address. The destination of the packet is known as the inside global address. When the same packet gets switched to the inside network, the source address is known as the outside local address and the destination of the packet is known as the inside local address.

This image provides an example.

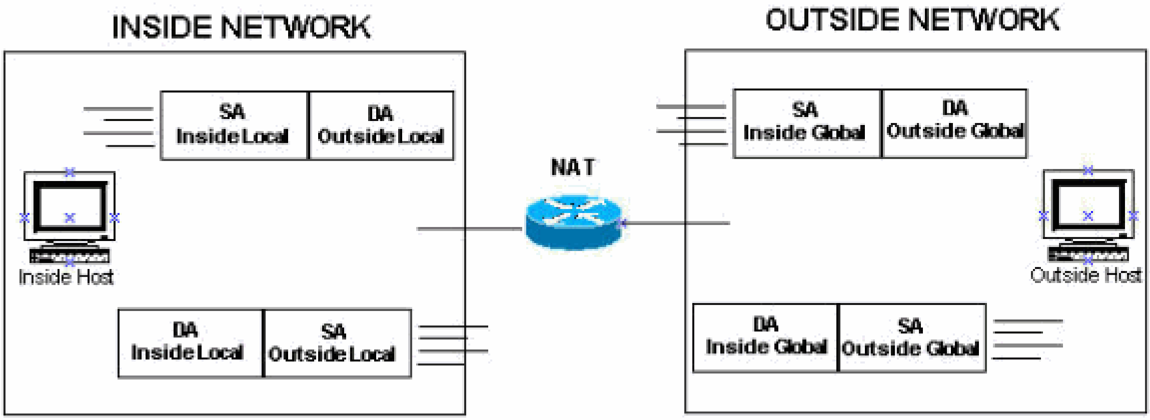

Examples

These sections examine these terms more closely and use this topology and examples.

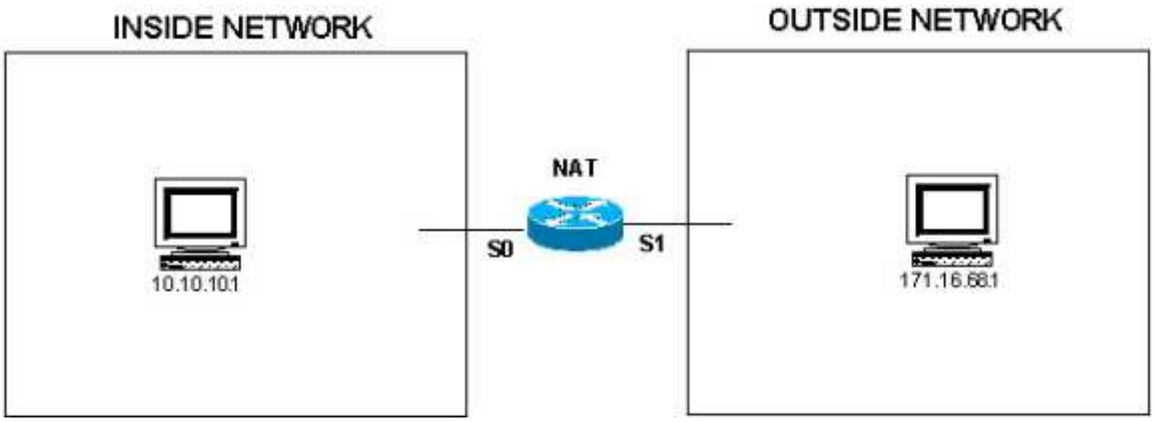

Define Inside Local and Inside Global Addresses

In this configuration, when the NAT router receives a packet on its inside interface with a source address of 10.10.10.1, the source address is translated to 171.16.68.5. This also means that when the NAT router receives a packet on its outside interface with a destination address of 171.16.68.5, the destination address is translated to 10.10.10.1.

```
ip nat inside source static 10.10.10.1 171.16.68.5

!--- Inside host is known by the outside host as 171.16.68.5.


interface s 0
ip nat inside

interface s 1
ip nat outside
```

You can issue the **show ip nat translations** command in order to verify the NAT translations in the router. In the ideal condition, the output of the **show ip nat translations** command is as shown here:

```
Router#show ip nat translations

Pro      Inside global       Inside local        Outside local
Outside global
---      171.16.68.5         10.10.10.1             ---
---
```

When the packet moves from the inside network to the outside network, the output of the **show ip nat translations** command is as shown here:

```
Router#show ip nat translations

Pro        Inside global          Inside local          Outside local
Outside global
icmp       171.16.68.5:15         10.10.10.1:15         171.16.68.1:15
171.16.68.1:15
---        171.16.68.5            10.10.10.1               ---
---
```

Note: In this output of the NAT translations, the protocol entry shows ICMP because Ping is used to validate the entries. The Outside Local and Outside Global entries will have the same IP address of the Outside host, which is 171.16.68.1.

The local addresses are addresses that appear on the inside cloud. Global addresses are addresses that appear on the outside cloud. Because of the way NAT is configured, the inside addresses are the only addresses that are translated. Therefore, the inside local address is different from the inside global address.

This is what the packets look like when they are on the inside network and on the outside network.

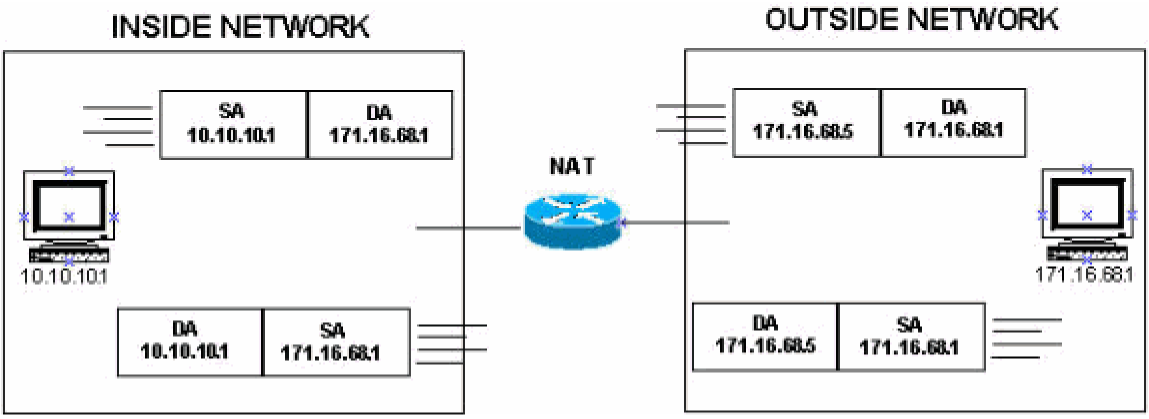

Define Outside Local and Outside Global Addresses

In this configuration, when the NAT router receives a packet on its outside interface with a source address of 171.16.68.1, the source address is translated to 10.10.10.5. This also means that if the NAT router receives a packet on its inside interface with a destination address of 10.10.10.5, the destination address is translated to 171.16.68.1.

```
ip nat outside source static 171.16.68.1 10.10.10.5

!--- Outside host is known to the inside host as 10.10.10.5.


interface s 0
ip nat inside

interface s 1
ip nat outside
```

In the ideal condition, the output of the **show ip nat translations** command is as shown here:

```
Router#show ip nat translations

Pro     Inside global           Inside local           Outside local
Outside global
      --- ---                        ---                  10.10.10.5
171.16.68.1
```

When the packet moves from the outside network to the inside network, the output of the **show ip nat translations** command is as shown here:

```
Router#show ip nat translations
```

```
Pro          Inside global        Inside local          Outside local
Outside global
                 --- ---              ---                 10.10.10.5
171.16.68.1
icmp         10.10.10.1:37        10.10.10.1:37         10.10.10.5:37
171.16.68.1:37
```

Note: The Inside Global and Inside Local entries will have the same IP address of the Inside host, which is 10.10.10.1.

The local addresses are addresses that appear on the inside cloud. Global addresses are addresses that appear on the outside cloud. In this example, because of the way NAT is configured, only the outside addresses get translated. Therefore, the outside local address is different from the outside global address.

This is what the packets look like when they are on the inside network and on the outside network.

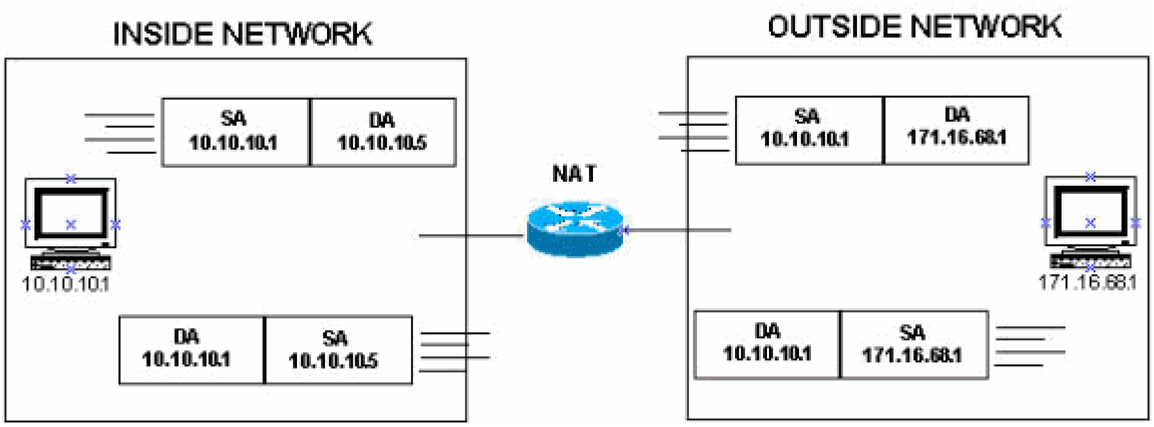

Define All Local and Global Addresses

In the this configuration, when the NAT router receives a packet on its inside interface with a source address of 10.10.10.1, the source address is translated to 171.16.68.5. When the NAT router receives a packet on its outside interface with a source address of 171.16.68.1, the source address is translated to 10.10.10.5.

This also means that when the NAT router receives a packet on its outside interface with a destination address of 171.16.68.5, the destination address is translated to 10.10.10.1. Also, when the NAT router receives a packet on its inside interface with a destination address of 10.10.10.5, the destination address is translated to 171.16.68.1.

```
ip nat inside source static 10.10.10.1 171.16.68.5

!--- Inside host is known to the outside host as 171.16.68.5.


ip nat outside source static 171.16.68.1 10.10.10.5

!--- Outside host is known to the inside host as 10.10.10.5.


interface s 0
ip nat inside

interface s 1
ip nat outside
```

In the ideal condition, the output of the **show ip nat translations** command is as shown here:

```
Router#show ip nat translations

Pro   Inside global      Inside local         Outside local
Outside global
         --- ---           ---                 10.10.10.5
171.16.68.1
       171.16.68.5        10.10.10.1            ---
---
```

The local addresses are addresses that appear on the inside cloud, and the global addresses are addresses that appear on the outside cloud. Because of how NAT is configured in this case, both the inside addresses and the outside addresses are translated. Therefore, the inside local addresses are different from the inside global addresses and the outside local addresses are different from the outside global addresses.

When the packet transfer is initiated from both the sides, the output of the **show ip nat translations** command is as shown here:

```
Router#show ip nat translations

Pro Inside global      Inside local       Outside local      Outside
global
---      ---                ---           10.10.10.5
171.16.68.1
icmp 10.10.10.1:4      10.10.10.1:4       10.10.10.5:4
171.16.68.1:4
icmp 171.16.68.5:39    10.10.10.1:39      171.16.68.1:39
171.16.68.1:39
---  171.16.68.5       10.10.10.1           ---                 ---
```

This is what the packets look like when they are on the inside network and on the outside network.

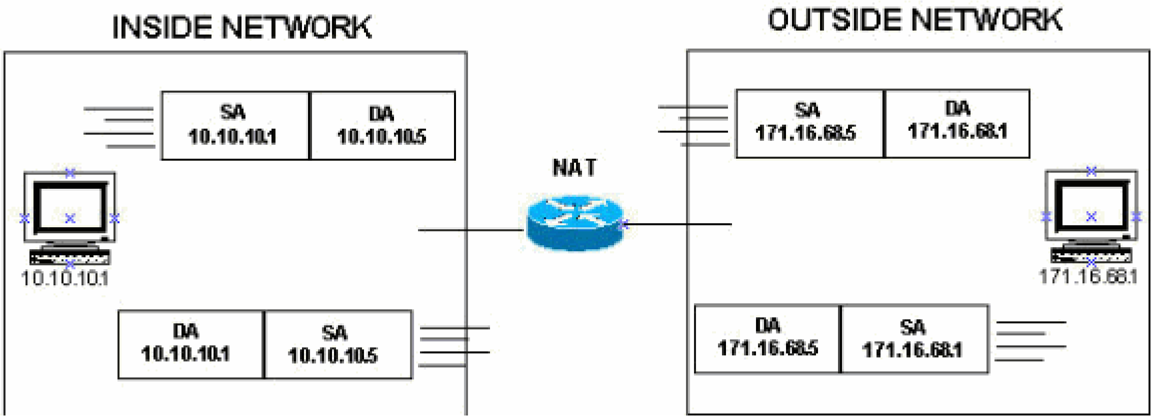

In summary, the terms local and global are actually very straight forward when you think of them in terms of where they appear in the network. Local addresses appear on the inside portion of the network while global addresses appear on the outside portion of the network.

Related Information

- Configuring Network Address Translation: Getting Started
- NAT Support Page
- IP Routing Support Page
- Technical Support & Documentation - Cisco Systems

Internet Connection Sharing with Windows XP

Published: July 1, 2002

By Sharon Crawford, Windows XP Expert Zone Community Columnist

Editor's Note: Past articles by members of the online community are archived for your use. The information may become outdated as technology changes. For the most current information, please search the Web site or post a question in the newsgroups.

Organizations with networks that include a connection to the Internet have hubs and routers to enable all the users to share the connection. For security, they'll also have at least one computer dedicated to being a proxy server or firewall of some kind. This is not a trivial matter to set up. Before Windows XP, you had to do something very similar at home to share a connection and make it secure while doing so. Most people didn't bother because their dial-up connections were made intermittently and they didn't feel that they were online long enough at any one time to be at risk.

Related Links

- Ask questions or discuss this topic in the Windows XP Expert Zone Newsgroups
- Windows XP Networking and the Web Newsgroup
- Set Up and Use Internet Connection Sharing
- Previous Columns by Sharon Crawford

However, now there are many home users with always-on Internet connections such as DSL and cable modems. Without protection, a computer that's always connected to the Internet is a sitting duck for malicious hackers.

Windows XP Professional and Home Edition come with two great services, **Internet Connection Sharing** (ICS) and **Internet Connection Firewall** (ICF). ICS and ICF allow a home user to share an always-on connection with security and without buying a license for expensive firewall software. Even a dial-up link can benefit from ICS and ICF, providing network address translation, addressing, and name resolution services for all the computers on your network plus security.

Setting Up Internet Connection Sharing

You will need to designate a Windows XP computer as the host. This computer must have two network adapters, one for your internal network and one for the Internet connection. Before attempting to enable ICS, verify that the host computer has a working connection to the Internet through the network card connected to the cable modem or DSL line, or on the network connection associated with the modem. The easiest way to enable ICS is to use the Network Setup Wizard, by following these steps:

1. Click **Start**, point to **All Programs**, point to Accessories, point to Communications, and then click **Network Setup Wizard**.
2. Click **Next** until you see the **Select a connection method** screen.
3. Click **This computer connects directly to the Internet**, and complete the wizard to install ICS.

This method has several advantages in that the wizard automatically detects the connection to the Internet, configures Internet Connection Firewall (ICF), bridges multiple network adapters connected to your home network and creates a log of information about the configuration named nsw.log in the Windows folder.

Turning on ICS manually is almost as easy as using the wizard except that you need to create the bridge for multiple network cards before enabling ICS. (See an earlier column, Building Network Bridges for more information on how to use the bridging capability in Windows XP.) Then take these steps:

1. In **Control Panel**, click **Network and Internet Connections** and then click **Network Connections**.
2. Click the local area network (LAN) connection or the dial-up networking connection that you want to share (that is, the one that connects to the Internet), and then, under Network Tasks, click **Change settings of this connection**.
3. Disable Client for Microsoft Networks and File and Print Sharing for Microsoft Networks by clearing the check boxes shown in Figure 1. This step is extremely **important**. Never leave these items enabled for any network card that is directly connected to the Internet (see sitting duck, above).

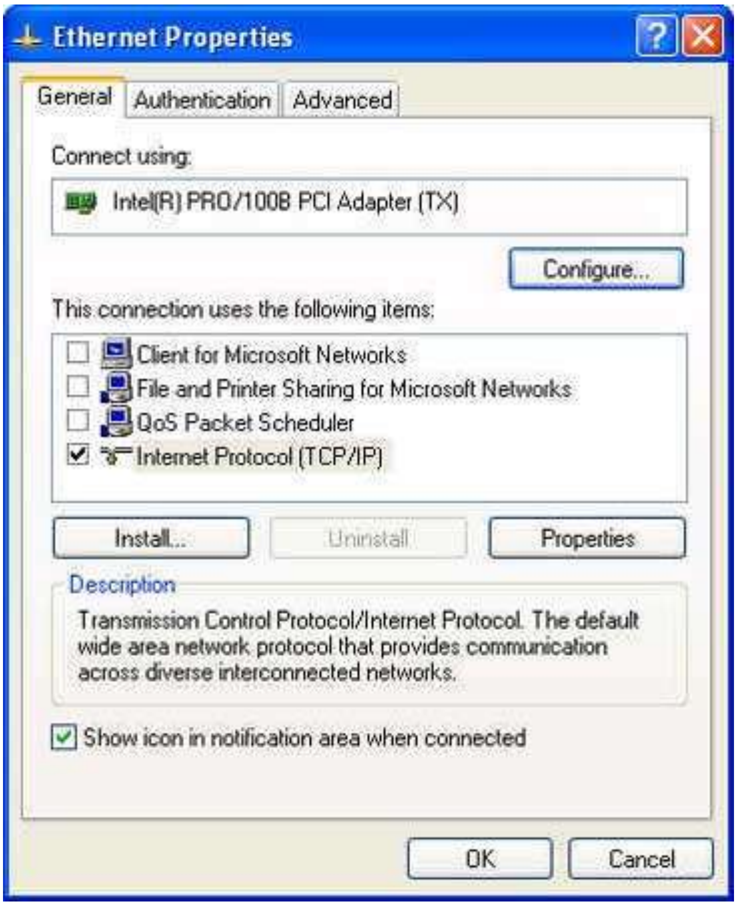

Connect using:

Intel(R) PRO/100B PCI Adapter (TX)

This connection uses the following items:

- □ Client for Microsoft Networks
- □ File and Printer Sharing for Microsoft Networks
- □ QoS Packet Scheduler
- ☑ Internet Protocol (TCP/IP)

Description

Transmission Control Protocol/Internet Protocol. The default wide area network protocol that provides communication across diverse interconnected networks.

☑ Show icon in notification area when connected

Figure 1

4. Click the **Advanced** tab, and select the **Allow other network users to connect through this computer's Internet connection** check box.

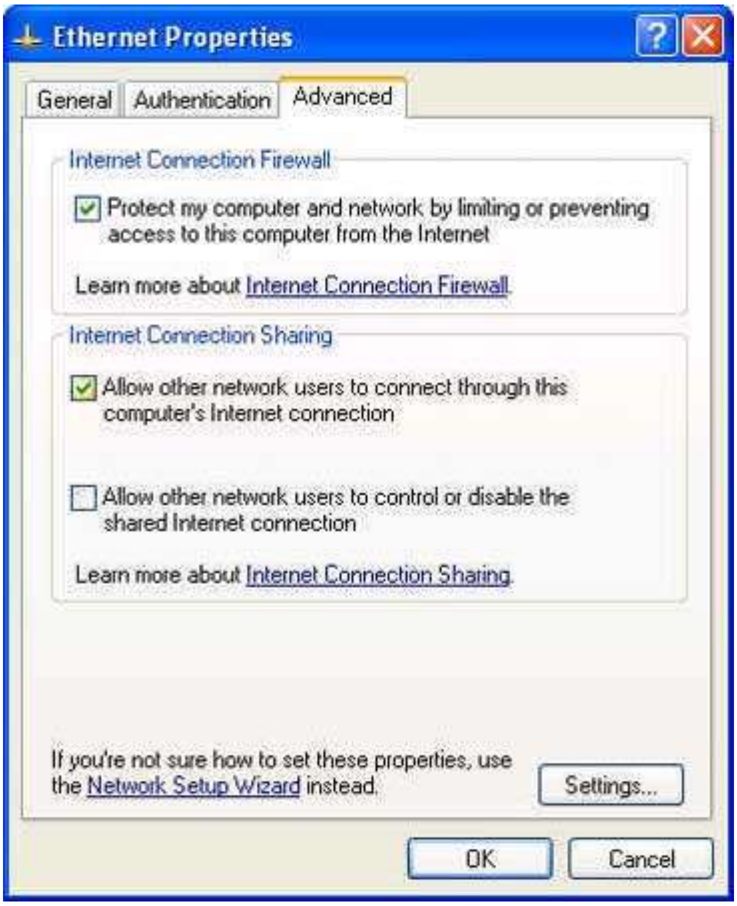

Internet Connection Firewall

☑ Protect my computer and network by limiting or preventing access to this computer from the Internet

Learn more about Internet Connection Firewall.

Internet Connection Sharing

☑ Allow other network users to connect through this computer's Internet connection

☐ Allow other network users to control or disable the shared Internet connection

Learn more about Internet Connection Sharing.

If you're not sure how to set these properties, use the Network Setup Wizard instead.

Figure 2

5. You can enable or disable the allowing of other users to control the connection—users don't need to be able to control the connection to use it.

6. Under Internet Connection Firewall, select the **Protect my computer and network by limiting or preventing access to this computer from the Internet** check box for this network card, unless you have another firewall between the computer and the Internet. This is very **important**.

7. Click **OK**, and Internet Connection Sharing will be enabled.

   Note: You must have administrative rights to enable ICS. After enabling ICS, verify that Internet connectivity is still functional on the host computer before testing the client computers. Remember to leave the host computer on all the time or turn it on before the other networked computers, so the client computers can request an IP address from the host.

↑Top of page

Troubleshooting ICS

If you have a problem with ICS, the best place to start is the Internet Connection Sharing Troubleshooter. You start the Troubleshooter with the following steps:

1. Click **Start**, and then click **Help and Support**.

2. Under **Pick a Help Topic**, click **Fixing a problem**.
3. In the left pane, click **Networking problems**.
4. In the right pane, click **Internet Connection Sharing Troubleshooter** and follow the instructions.

The Troubleshooter can address problems such as not being able to receive e-mail on an ICS client, the client or host computer fails to dial out or dials out without notifying you, you're unable to browse the Internet from a client or host computer, or your DSL or cable modem connection is slow. However, if the Troubleshooter leaves you troubled, here are some other common problems and their solutions.

⇡Top of page

ICS Not Enabled

If you're configuring ICS manually, be sure that the internal network adapter on the host computer doesn't have Internet Connection Firewall enabled. If ICF is enabled, you'll have to disable it before configuring ICS on the external adapter. Or take the easy way and run the Network Setup Wizard, which will automatically disable ICF on home networking adapters.

Check the IP address on the external adapter to verify that it is obtaining an IP address from your ISP. Similarly, check the IP address on the internal network adapter to verify that it is 192.168.0.1. If it's not, disable ICS, and then make sure the internal adapter is configured to use DCHP. Then re-enable ICS.

Internet Connection Sharing (ICS) automates the IP numbering task for the ICS clients on your network with the Dynamic Host Configuration Protocol (DHCP) service. The DHCP service enables the ICS host computer to assign IP addresses to its clients automatically. By default, when ICS is installed, the DHCP service begins supplying addresses to computers on the network.

⇡Top of page

Cannot Print to a Network Printer after Adding ICS

After you add Internet Connection Sharing (ICS), you discover that you can't print. This can happen because ICS uses a Class C subnet with an address range of 198.168.0.x. To solve the problem, give the printer an IP address to match the subnet of the client computers.

⇡Top of page

Computers on the Network Can't Connect to the Host

As part of the process of enabling ICS, the network adapter for the internal network on the host computer is set to a fixed IP address of 192.168.0.1 and a special DHCP server is enabled on that connection.

If computers on your network can't see the ICS host, it may be because they are not enabled to use DHCP. Check to see if DHCP is enabled on the client computer:

1. In **Control Panel**, click **Network and Internet Connections**, and then click **Network Connections**.
2. Right-click the connection icon, and then click **Properties**.
3. Highlight **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. On the **General** tab, if an IP address is specified, select the option **Obtain an IP address automatically**.

If a client computer has DHCP enabled and still can't see the host computer, try rebooting the client. Make sure that there are no other DHCP providers on the network, such as an Internet gateway device. Any such device should be on the outside segment of the network—between the host computer and the Internet, not between the host computer and the internal network.

If you use Windows XP at home or in a small business, and you have a topic you'd like to see covered in a future column, feel free to write me at: sharoncrawford@mvps.org. I'd be glad to receive ideas and suggestions.

Sharon Crawford is a former editor now engaged in writing books and magazine articles. Since 1993, she has written or co-written two dozen books on computer topics. Her books include Windows 2000 Pro: The Missing Manual, Windows 98: No Experience Required, *and* Windows 2000 Professional for Dummies (with Andy Rathbone).

How to Install Internet Connection Sharing (ICS)

Internet Connection Sharing (ICS) is part of Windows 98 Second Edition (or Windows 2000), so you'll need either one of these operating systems to be able to install ICS.

Notes before getting started:

- If any of your client computers are set to **Obtain an IP address automatically** (from a DHCP server), you should shut down those computers before you install ICS. That is to make sure that the IP information assigned by the old DHCP server doesn't interfere with the information assigned by the ICS DHCP server.
- If you already have another sharing application like Sygate or Wingate installed, uninstall it, before installing ICS. Some sharing programs (ICS included) take control of one or more of your Network adapters and/or Protocols, and having more then one program trying to control them can lead to conflicts.
- If you're going to share a cable modem or ASDL connection with ICS, you'll need two Network Interface Cards (NICs) in the computer that you install ICS on. (The ICS installer doesn't give you any alternative.)
- The Microsoft ATM adapter does not work with ICS.
- Hybrid services like DirecPC and one-way cable modems with telco modem returns do not work with ICS.

Installing ICS:

1. After you install Windows 98SE, select **Start** > **Settings** > **Control Panel**, and double-click **Add/Remove Programs**
2. Select the **Windows Setup** tab
3. Double-click the **Internet Tools** icon, and put a check-mark in the box for **Internet Connection Sharing**
4. Click **OK** to close the **Internet Tools** window, and click **OK** again to close the **Add/Remove Programs** window
5. ICS will launch its setup wizard to guide you through the process

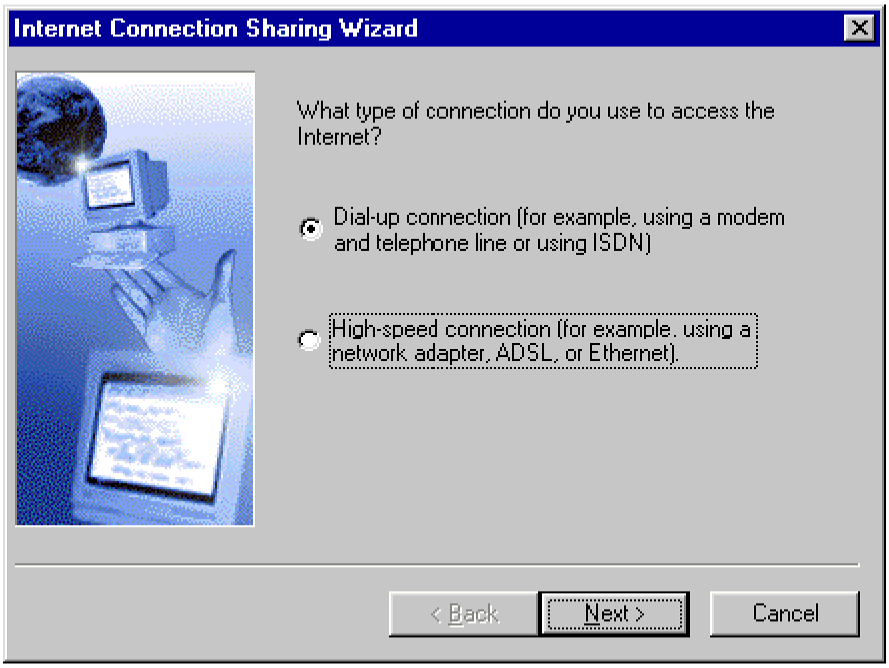

Internet Connection Sharing Wizard

What type of connection do you use to access the Internet?

● Dial-up connection (for example, using a modem and telephone line or using ISDN)

○ High-speed connection (for example, using a network adapter, ADSL, or Ethernet).

The first selection you need to make in the Wizard is for the type of connection you will be using. Make sure you select the right one.

If you are using a cable modem, you will be asked to select a Network adapter (NIC). Be sure you choose the adapter that is attached to your cable modem. If you don't choose the correct NIC, ICS won't install correctly and you'll probably have to remove and reinstall it and try again. You will see that the ICS Wizard numbers your NIC's #1 and #2, but here's no clue as to which is connected to the cable modem and which is connected to your LAN. You have to guess, and if you guess wrong, you get to go through the whole process again!

Next the Wizard will tell you to create a **Client Configuration Disk**.

The next window will finish the installation and reboot your system.

Setting Up NAT on Solaris Using IP Filter

Introduction

So, you've got several computers on your home or business network, and you'd like to be able to access the Internet from all of them, probably via a cable (or DSL) modem. Basically you have three options:

1. You connect all your machines and your cable modem to a hub, set them all up as DHCP clients (see this page for how to do this on Solaris), and go for it.
2. You set up one of your machines to do NAT (Network Address Translation), hiding the rest behind a firewall using RFC 1918 compliant addresses on your network.
3. You use one of those Netgear routers, or someting similar (e.g., those from Linksys), as your firewall, and let it perform NAT for you.

The last option is very popular, and is better than nothing, but you can't beat having your own dedicated firewall machine. The first method, as well as being insecure, lacks a certain *je ne sais quoi*, so I'll show you how to set up NAT using Darren Reed's IP Filter. If you want to use the first or last methods, you're on your own!

Hardware

In my experiments, I could only get NAT to work reliably when I had two physical interfaces (i.e., using two virtual interfaces, say `hme0` and `hme0:1`, didn't work). I used `hme1` to connect directly to my cable modem, and `hme0` as the connection to the rest of my network via a 100 baseT switch. `hme1` is under DHCP control per these instructions, and `hme0` was set up the conventional way, with the hostname in `/etc/hostname.hme0`, and the corresponding IP address in `/etc/hosts`.

Installing IP Filter

By far the best way to get IP Filter is install Solaris 10, which comes with Solaris IP Filter (which is based on IP Filter). For previous versions of Solaris, the best way to get IP Filter is to compile a copy of the latest source code, which can be downloaded from the IP Filter home page. As an alternative, I have a compiled version of the package here. This is IP Filter version 3.3.11, compiled on a Sun SPARCstation 20, running Solaris 2.6. I've also used it on a SPARCstation 2 running Solaris 7, but it is provided here without any support (I currently use the Solaris 10 version of IP Filter on a Sun Netra T1 105). You should probably download a more recent binary from Marauding Pirates.

Configuring IP Filter on Solaris 10

Once you've successfully installed IP Filter, you need to configure it. First of all, you need to make sure that your NAT box will forward IP packets (it's possible this ability was disabled for security reasons). As root, run this command:

```
    routeadm
```

If the "Current Configuration" column of the "IPv4 forwarding" row says "disabled", then you must enable it. You do this by running the following command (again, as root):

```
routeadm -u -e ipv4-forwarding
```

The `-e ipv4-forwarding` option causes IPv4 forwarding to be enabled, and the `-u` flag causes the change to be applied to the running system (in addition to changing the settings when the system is next rebooted).

When you're happy that IP forwarding is enabled, you need to set up your NAT rules. The file `/etc/ipf/ipnat.conf` contains the rules you want to use. This is the `ipnat.conf` file I use, bearing in mind that all of my machines have an IP address in the 192.168.0.1 to 192.168.0.254 range; you should change the addresses between "`hme1`" and the "->" to suit your needs (note also that I've specified `hme1`; put the name of your outbound interface here instead):

```
map hme1 192.168.0.0/24 -> 0/32 proxy port ftp ftp/tcp
map hme1 192.168.0.0/24 -> 0/32 portmap tcp/udp auto
map hme1 192.168.0.0/24 -> 0/32
```

The `0/32` stuff is some magic to tell IP Filter to use the address currently assigned to the interface - very useful in DHCP client environments!

The order of the rules is important; don't change them unless you know what you're doing, otherwise things will break! The first rule allows FTP access from all of your hosts. The second maps the source port numbers to a high range (10000 to 40000 by default), and the third rule maps all other TCP traffic.

Once you've set up your NAT rules, you need to enable packet filtering for the interface type you're using. This is done by uncommenting the appropriate line(s) in `/etc/ipf/pfil.ap`:

```
#le     -1      0       pfil
#qe     -1      0       pfil
hme     -1      0       pfil
```

When you're happy with your configuration, start the IP filter services:

```
svcadm restart network/pfil
svcadm restart ipfilter
```

The interfaces that you enabled packet filtering on by editing `/etc/ipf/pfil.ap` must be replumbed before you can use them. Here's how to do it, assuming your machine is set up like mine:

```
ifconfig hme1 unplumb
ifconfig hme1 plumb dhcp start
```

Another, perhaps easier, way is to simply reboot your machine. Although it smells like a typical Windoze "admin" kind of way of doing this, it does have the advantage of testing that your modifications will survive a reboot.

Assuming all is well, your firewall should now correctly handle NAT, even after a reboot. Assuming this is the case, enjoy! If this page has been useful to you, please consider buying a copy of my book, Solaris Systems Programming.

Configuring IP Filter for Previous Versions of Solaris

If you're using a version of Solaris prior to Solaris 10, and assuming you have Solaris 10-capable hardware, I don't know why you **wouldn't** use Solaris 10, here is the older version of these instructions. But really, you should upgrade to Solaris 10!

First of all, you need to make sure that your NAT box will forward IP packets (it's possible this ability was disabled for security reasons). As root, run this command:

```
ndd -get /dev/tcp ip_forwarding
```

If the result is "1", you're all set. Zero means that IP forwarding is not enabled. To enable it, delete the file `/etc/notrouter`, and possibly `/etc/defaultrouter` too. Create an empty `/etc/gateways` file, and IP forwarding will be enabled at the next reboot.

One caveat applies, though: if you're using NAT and DHCP on the same server (like I do), IP forwarding will not get enabled. So, I install this script as `/etc/init.d/ip_forwarding`, with a symbolic link to it from `/etc/rc2.d/S69ip_forwarding`. With this script in place, IP forwarding will be enabled even if you are using a DHCP client.

When you're happy that IP Filter is running, and IP forwarding is enabled, you need to set up your NAT rules. The file `/etc/opt/ipf/ipnat.conf` contains the rules you want to use. This is the `ipnat.conf` file I use, bearing in mind that all of my machines have an IP address in the 192.168.0.1 to 192.168.0.254 range; you should change the addresses between "`hme1`" and the "->" to suit your needs (note also that I've specified `hme1`; put the name of your outbound interface here instead):

```
map hme1 192.168.0.0/24 -> 0/32 proxy port ftp ftp/tcp
map hme1 192.168.0.0/24 -> 0/32 portmap tcp/udp auto
map hme1 192.168.0.0/24 -> 0/32
```

The `0/32` stuff is some magic to tell IP Filter to use the address currently assigned to the interface - very useful in DHCP client environments!

The order of the rules is important; don't change them unless you know what you're doing, otherwise things will break! The first rule allows FTP access from all of your hosts. The second maps the source port numbers to a high range (10000 to 40000 by default), and the third rule maps all other TCP traffic.

Use `/etc/init.d/ipfboot stop` and `/etc/init.d/ipfboot start` to test your configuration, and when you're happy that all is working well, reboot. This will make sure that everything still works as expected, even after a reboot.

That's about it - enjoy! If this page has been useful to you, please consider buying a copy of my book, Solaris Systems Programming.

NAT Tools and Settings

Updated: July 31, 2009

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

In this section

- NAT Tools
- NAT Registry Entries
- Related Information

You can use the network address translation (NAT) tools and registry settings described here to enable, configure, and manage Routing and Remote Access NAT on a computer running Windows Server 2003.

NAT Tools

The following tools are associated with the NAT routing protocol component provided by the Microsoft Windows Server 2003 Routing and Remote Access service:

- Routing and Remote Access snap-in
- Network Connections
- Netsh command-line tools for Routing and Remote Access NAT

Graphical User Interface Tools

The graphical user interface tools used to install and configure Routing and Remote Access NAT include Network Connections, used to configure TCP/IP properties for NAT clients to provide the client computers access to the Internet, and the Routing and Remote Access snap-in, used to install and configure Routing and Remote Access NAT on a server.

Network Connections

Category

Network Connections is included with the versions of the Windows operating systems as described in the next paragraph, “Version compatibility.”

Version compatibility

You can run Network Connections on any computer running Windows Server 2003, Windows XP, Windows 2000, or Windows NT version 4.0. However, Windows NT 4.0 does not support Routing and Remote Access NAT.

Network Connections options for NAT clients

You use the Network Connections tool to configure TCP/IP properties for NAT clients so that the client computers can use a NAT-enabled router to gain access to the Internet (or other public network).

The properties page for TCP/IP on client computers can be used for a variety of purposes. The following table lists which TCP/IP options are used to enable NAT clients to interact with a NAT-enabled router on a private network.

TCP/IP Properties Options Used for a NAT Client

Page	Tab	Option
TCP/IP Properties	General	The following options on the adapter of the NAT client are used to enable the NAT client to access the NAT-enabled router: 1. Under **Use the following IP address**, the values needed for following options are specific to a NAT client: • **IP address.** The private IPv4 address and subnet mask of the client. • **Subnet mask.** The private subnet mask of the client. • **Default gateway.** The private IPv4 address of the NAT-enabled router. 2. Under **Use the following DNS server addresses**, the value needed for the following option is specific to a NAT client: **Preferred DNS server.** The private IPv4 address of the NAT-enabled router, which is acting as a DNS proxy.
TCP/IP Properties	Alternate Configuration	No NAT-related configuration is needed on this tab.
Advanced TCP/IP Settings	IP Settings	**Default gateways.** A gateway using the private address of the NAT-enabled router (if appropriate). **Caution:** This option is appropriate to use only if the client computer receives its IP address from a DHCP server (that is, this option is not appropriate if the client obtains its IP address from the NAT DHCP allocator).

Routing and Remote Access Snap-in

Category

The Routing and Remote Access snap-in is used with the Routing and Remote Access service, which is included with Windows Server 2003 and Windows 2000 Server. Routing and Remote Access is disabled by default. You can use the Routing and Remote Access snap-in, under **Administrative Tools**, to enable and configure the Routing and Remote Access service, including the optional NAT routing protocol component.

Version compatibility

The Routing and Remote Access snap-in is provided by the Routing and Remote Access service on computers running Windows Server 2003 and Windows 2000 Server. For Windows NT 4.0, the Routing and Remote Access

Service (RRAS) and its snap-in are available as a separate download from the Microsoft Windows NT Server Routing and Remote Access Service Download page at http://go.microsoft.com/fwlink/?LinkId=22441. However, Windows NT 4.0 RRAS does not include NAT.

The following sections briefly summarize NAT-related tasks for which you can use the Routing and Remote Access snap-in. (In addition to NAT, the Routing and Remote Access snap-in is also used to configure LAN routing, dial-up or VPN remote access connections, and site-to-site connections between geographically remote networks.)

- Enabling NAT while running the Routing and Remote Access wizard
- Enabling both VPN and NAT while running the Routing and Remote Access wizard

Enabling NAT while running the Routing and Remote Access Wizard

If the Routing and Remote Access service is not yet enabled on a computer running Windows Server 2003, you can enable NAT when you run the Routing and Remote Access Wizard by selecting the **Network Address Translation (NAT)** option. The wizard also lets you choose to enable the Basic Firewall feature.

Choosing **Network Address Translation (NAT)** when you run the wizard establishes the following:

- Configures the IP address of the private network interface (the LAN card that connects to the private network segment).
- Configures the public interface. If the connection is a non-permanent connection (such as a dial-up modem), the wizard creates a demand-dial interface to the ISP and creates a default static route that uses the Internet interface. (If the connection is permanent, such as DDS, T-Carrier, Frame Relay, permanent ISDN, xDSL, or cable modem, the wizard does not create a demand-dial interface or static route for the interface.)
- Adds the NAT routing protocol component.
- Adds Internet and private network interfaces to the NAT routing protocol component.
- If you chose the option to enable **Basic Firewall** while running the wizard, the wizard configures a basic stateful firewall on the public interface connected to the Internet.

  Note

  - If the network already has a firewall and you do not select the **Basic Firewall** option while running the wizard, the Routing and Remote Access snap-in entry for the NAT routing protocol component (under **IP Routing** in the console tree) displays as **NAT/Basic Firewall**. The name “NAT/Basic Firewall” does not indicate whether Basic Firewall is configured.
  - You can confirm whether Basic Firewall is configured by using the **NAT/Basic Firewall** tab on the properties page of the public (Internet-connected) interface.

Enabling both VPN and NAT while running the Routing and Remote Access Wizard

If the Routing and Remote Access service is not yet enabled on a computer running Windows Server 2003, you can configure the server both to provide NAT for the private network and also to accept VPN connections. You can do so when you run the Routing and Remote Access Wizard by selecting the **Virtual Private Network (VPN) access and NAT** option when the wizard begins.

Choosing **Virtual Private Network (VPN) access and NAT** specifies that computers on the Internet cannot determine the IP addresses of any computer on the private network, yet allows VPN clients to connect to computers on the private network.

Routing and Remote Access snap-in options for NAT-enabled routers

If the Routing and Remote Access service is already enabled on a server, or if you installed NAT by using the Routing and Remote Access Wizard and want to modify the NAT configuration, you can use the tools provided by the Routing and Remote Access snap-in to enable and configure, or modify, Routing and Remote Access NAT.

The Routing and Remote Access snap-in can be used for a variety of purposes unrelated to Routing and Remote Access NAT. The following table lists which options under the **General** and **NAT/Basic Firewall** nodes in the Routing and Remote Access snap-in are used for NAT-related tasks and describes the location in the Routing and Remote Access snap-in used for each task.

NAT-related Options in the Routing and Remote Access Console Tree

Node	Task
General	Adding network address translation: • Under the server name for the server to be configured as the NAT-enabled router, expand **IP Routing**, right-click **General**, select **New Routing Protocol**, and then choose **NAT/Basic Firewall**.
NAT/Basic Firewall	Adding and configuring public or private interfaces for the NAT routing protocol component: • Right-click **NAT/Basic Firewall**, and then click **New Interface** to add an internal interface to connect to the private network or to add a public interface to connect to the Internet. **Note:** You do not need to manually configure public or private interfaces for the NAT component if you used the Routing and Remote Access Setup wizard to configure NAT. Viewing the NAT mapping table: • Click **NAT/Basic Firewall**, right-click the public interface in the details pane, and then click **Show Mappings**. Viewing DHCP allocator information: • Right-click **NAT/Basic Firewall**, and then select **Show DHCP Allocator Information** to display the number of instances for each of the following: • Messages ignored • DECLINE messages received • DISCOVER messages received • INFORM messages received

	• RELEASE messages received • REQUEST messages received • ACK messages sent • BOOTP replies sent • NAK messages sent • OFFER messages sent Viewing DNS proxy information: • Right-click **NAT/Basic Firewall**, and then select **Show DNS Proxy Information** to display the number of instances for each of the following: • Messages ignored • Queries received • Responses received • Queries sent • Responses sent

The following table describes how each tab on the **NAT/Basic Firewall Properties** page in the Routing and Remote Access snap-in is used for NAT-related tasks.

NAT-related Options on the NAT/Basic Firewall Properties Page

Tab	Task
General	Specifying the level of errors and warnings to be logged in the System Log in Event Viewer: • **Log errors only.** Specifies that only errors are logged in the System Log in Event Viewer. • **Log errors and messages.** Specifies that both errors and warnings are logged in the System Log in Event Viewer. • **Log the maximum amount of information.** Specifies that the maximum amount of information is logged in the System Log in Event Viewer. • **Disable event logging.** Specifies that no events are logged in the System Log

	in Event Viewer.
Translation	Specifying the number of minutes that a dynamic mapping for a TCP session or for a UDP message remains in the NAT Mapping Table.
Address Assignment	Configuring the DHCP allocator feature: 1. Specify whether the NAT-enabled router will provide DHCP-based address assignment to DHCP clients on the private network. 2. Specify both the private address range and any exclusions; that is, specify any addresses within the specified range of addresses that should not be assigned to DHCP clients on the private network because they are already in use. If multiple routed subnets are configured, you must use a DHCP server rather than the DHCP allocator.
Name Resolution	Configure the DNS proxy feature: 1. Specify whether the NAT-enabled router relays DNS name resolution requests from hosts on the private network to the configured DNS server for the NAT-enabled router. 2. Specify whether a connection is attempted by using the selected demand-dial interface when a DNS name resolution request is received by a host on the private network.

The following table describes how each tab on the properties page of the public (Internet-connected) interface in the details pane of the Routing and Remote Access snap-in is used for NAT-related tasks.

NAT-related Options on the Public Interface Properties Page

Tab	Task
NAT/Basic Firewall	Configuring NAT: • Select **Enable NAT on this interface** to enable the router to send data to and receive data from the Internet over this interface. Configuring Basic Firewall: • Select **Enable a basic firewall on this interface** to protect computers on the private network from unsolicited Internet traffic. Configuring static packet filters: • Under **Static packet filters**, select **Inbound Filters** or **Outbound** Filters to establish inbound and outbound packet filters on the public interface to restrict traffic

	based on packet attributes such as IP address or protocol. For example, you can use this option to configure filters for PPTP or L2TP/IPSec VPN connections, as described in the Windows Server 2003 Deployment Guide section about “Configuring Packet Filters for a VPN Server” in Deploying a VPN Remote Access Server Solution.
Address Pool	Configuring one or more IP address ranges: • Select **Add** to configure an IP address pool on the public interface. You cannot configure an IP reservation until you configure at least one IP address pool. Configuring an IP reservation to allow incoming traffic to a computer (such as a Web server) on the private network: • Under **Reserve public addresses**, click **Reservations** to configure an IP reservation by specifying a public IP address from one of the configured ranges. This ensures that the reserved address cannot be used for address translation.
Services and Ports	Configuring a static mapping for the services on your network to which you want to provide access for Internet users: • Select one of more from the list of services provided (such as **FTP Server**, **Web Server (HTTP)**, or **Remote Desktop**). -or- • Click **Add** to configure a service that is not on the list to which you want to provide access for Internet users. **Note:** Configuring a service on this tab creates a static entry in the NAT Mapping Table and creates exceptions in the Basic Firewall that allow the specified incoming traffic.
ICMP	Configuring Internet Control Message Protocol (ICMP) options: • Select the requests for error and status information (ICMP messages) from the Internet to which this computer will respond.

The following table describes how the tab on the properties page of the **Internal** interface in the details pane of the Routing and Remote Access snap-in is used for NAT-related tasks.

NAT-related Options on the Internal Properties Page

Tab	Task
NAT/Basic Firewall	Configuring static packet filters: • Under **Static packet filters**, select **Inbound Filters** or **OutboundFilters** to establish inbound and outbound packet filters on the private interface to restrict

	traffic based on packet attributes such as IP address or protocol **Note:** If the NAT/Basic Firewall component is enabled, the option **Private interface connected to private network** is selected on this tab by default and cannot be unselected.

Netsh Command-Line Tools for Routing and Remote Access NAT

Netsh provides several sets of commands (also known as *contexts*) for performing a wide range of network configuration tasks. The Netsh Routing IP NAT commandsprovide the Netsh context for Routing and Remote Access NAT.

Netsh.exe: Netsh Routing IP NAT Commands

Category

The Netsh Routing IP NAT commands, a subset of the Netsh command-line toolset, are included with the Windows Server 2003 operating system.

Version compatibility

The Netsh Routing IP NAT commands are compatible with Windows Server 2003. Netsh commands were first introduced in Windows 2000 Server and were expanded to include additional commands, including commands to manage NAT, in Windows Server 2003.

The Netsh commands are designed to help network administrators manage a TCP/IP network. You can use the Netsh command-line set of tools to locally or remotely display or modify the configuration of services or protocols on Windows–based computers. The Netsh command-line interface is scriptable, which lets you perform batch configurations or network administration from a centralized location. In addition to the Netsh Routing IP NAT commands that are designed specifically for Routing and Remote Access NAT, NAT also inherits commands from the Netsh Routing context and the Netsh Routing IP context.

The following table contains a brief description of the commands available in the Netsh Routing IP NAT context.

Commands Available in the Netsh Routing IP NAT Context

NAT Context Command	Description
? or help	When typed at a **netsh routing ip nat>** prompt, either **?** or **help** displays a complete list of all commands in the Netsh Routing IP NAT context, including all commands inherited from the global Netsh context as well as commands inherited from the Netsh Routing and Netsh Routing IP subcontexts. When typed at a **netsh routing ip nat>** prompt, a command name followed by **?** (such as **show ?**) displays information about that command.
add addressmapping	Adds an IP address mapping to the NAT address pool for the specified interface.
add addressrange	Adds an address range to the NAT address pool for the specified interface.
add ftp	Enables the NAT proxy for FTP (supports FTP traffic across a NAT).
add h323	Enables the NAT proxy for H.323 (supports NetMeeting calls across a NAT).

add interface	Configures NAT on the specified interface.
add portmapping	Adds a protocol port mapping for either the TCP or the UDP protocol type on the NAT interface.
delete addressmapping	Deletes an address mapping from the NAT address pool for the specified interface.
delete addressrange	Deletes an address range from the NAT address pool for the specified interface.
delete ftp	Disables the NAT proxy for FTP.
delete h323	Disables the NAT proxy for H323.
delete interface	Removes NAT from the specified interface.
delete portmapping	Deletes a protocol port mapping for either the TCP or the UDP protocol type from the specified NAT-enabled interface.
set global	Sets the following global parameters for NAT: • The timeout value, in minutes, for TCP mappings. • The timeout value, in minutes, for UDP mappings. • Which events should be logged. The **none** parameter specifies that no events related to NAT should be logged. The **error** parameter specifies that only errors related to NAT should be logged. The **warn** parameter specifies that only warnings related to NAT should be logged. The **info** parameter specifies that all events related to NAT should be logged.
set interface	Configures NAT parameters for the specified interface.
show global	Displays NAT global configuration. That is, it displays the current defaults for the following: • TCP timeout (in minutes) • UDP timeout (in minutes) • Logging level (such as errors only)
show interface	Displays NAT configuration for the specified interface.

For more information about Netsh, see "Command-Line Reference for Windows Server 2003, Standard Edition" in the Tools and Settings Collection.

NAT Registry Entries

The following registry entry, associated with Routing and Remote Access NAT, is the only registry entry that an administrator might want to modify by using the registry editor.

The information here is provided as a reference for use in troubleshooting or verifying that the required settings are applied. It is recommended that you do not directly edit the registry unless there is no other alternative. Modifications to the registry are not validated by the registry editor or by Windows before they are applied, and as a result, incorrect values can be stored. This can result in unrecoverable errors in the system. When possible, use Group Policy or other Windows tools, such as Microsoft Management Console (MMC), to accomplish tasks rather than editing the registry directly. If you must edit the registry, use extreme caution.

AllowInboundNonUnicastTraffic
Registry path
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IpNat\Parameters\

Version
Windows Server 2003, Windows XP SP1, or later.

If Routing and Remote Access NAT has Basic Firewall configured, the firewall always accepts broadcast and multicast packets and passes them to the NAT component. However, on a computer running the Windows Server 2003, Windows XP SP1, or later operating system, the following registry key for NAT is set by default to drop all inbound broadcast and multicast packets. If you need to change this default behavior, add the following registry key and set it to **1**. Setting the key to **1** allows broadcast and multicast packets to cross Basic Firewall:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IpNat\Parameters\AllowInboundNonUnicastTraffic

By default, **AllowInboundNonUnicastTraffic** is set to **0**, which blocks inbound unicast traffic.

For more information about this registry entry, see the Registry Reference for Windows Server 2003.

Related Information

The following resources contain additional information that is relevant to this section:

- "Command-Line Reference for Windows Server 2003, Standard Edition" in the Tools and Settings Collection
- "Registry Reference for Windows Server 2003" in the Tools and Settings Collection

Certificate Services

Installing and configuring a certification authority

Updated: January 21, 2005

Applies To: Windows Server 2003, Windows Server 2003 R2, Windows Server 2003 with SP1, Windows Server 2003 with SP2

Installing and configuring a certification authority

This section will give you additional information you need to install and configure a certification authority (CA).

Planning the installation of a certification authority

Before installing Certificate Services, you should plan the deployment of certification authorities (CAs) and Deploying a Public Key Infrastructure in your organization.

For more information about deploying a PKI, see Certificates Resources in public key infrastructure and Microsoft Windows Deployment and Resource Kits.

Ways to install Certificate Services to create a certification authority

If you are installing Certificate Services, there are a number of situations in which you may want to set up a CA. (The most typical is listed first here.)

- **After base setup has completed**:

  To install Certificate Services on a server that already has Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; or Windows Server 2003, Datacenter Edition installed, open **Add or Remove Programs** in **Control Panel**. (To open a Control Panel item, click **Start**, click **Control Panel**, and then double-click the appropriate icon.) After you select Certificate Services for installation, the Certificate Services Installation Wizard guides you through the installation process.

  For the procedures to set up a certification authority (CA) as part of the base setup, see Set Up a Certification Authority.

- **As part of** Windows Server 2003 **family base setup**:

  Although Certificate Services is a service and is included with the Windows Server 2003 family, it is not installed as part of the initial installation process by default. (You do not want every file server to be a CA.)

  To install Certificate Services during the initial base installation of the Windows Server 2003 family, you must select it from the optional components list that is displayed during setup. Certificate Services will not actually be installed until you log on to the server after setup has completed. Then, a message will prompt you to complete the setup of the CA.

  This feature is not included on computers running the Microsoft® Windows Server® 2003, Web Edition, operating system. For more information, see Overview of Windows Server 2003, Web Edition.

Setup options and information

When you set up Certificate Services, you will be prompted for the following information:

Certification authority type selection

During the Certificate Services installation, you can choose to set up any of the following types of certification authority (CA):

Certification Authority Type	Description
Enterprise root CA	An enterprise root CA is a top-level CA in a certification hierarchy. An enterprise root CA requires the Active Directory directory service. It self-signs its own CA certificate and uses Group Policy to publish that certificate to the Trusted Root Certification Authorities store of all servers and workstations in the domain. Normally, an enterprise root CA does not directly provide resources for user and computer certificates, but is the foundation for a certificate hierarchy. For more information, see Enterprise certification

	authorities
Enterprise subordinate CA	An enterprise subordinate CA must obtain its CA certificate from another CA. An enterprise subordinate CA requires Active Directory. You use enterprise subordinate CAs when you want to take advantage of Active Directory, certificate templates, and smart card logon to Windows XP and computers running Windows Server 2003 family operating systems
Stand-alone root CA	A stand-alone root CA is a top-level CA in a certification hierarchy. The stand-alone root CA may or may not be a member of a domain and, therefore, does not require Active Directory. However, it will use Active Directory if it exists for publishing certificates and certificate revocation lists. Since a stand-alone root CA does not require Active Directory, it can easily be disconnected from the network and placed in a secure area, which is useful when creating a secure offline root CA. For more information, see Stand-alone certification authorities
Stand-alone subordinate CA	A stand-alone subordinate CA must obtain its CA certificate from another CA. The stand-alone subordinate CA may or may not be a member of a domain and, therefore, does not require Active Directory. However, it will use Active Directory if it exists for publishing certificates and certificate revocation lists. It must obtain its CA certificate from another CA.

Note

- Certificate Services should not be installed on any node in a server cluster as this configuration may not allow Certificate Services to run properly and is therefore not supported.

Installing Web enrollment support

You may want to install a certification authority infrastructure within a defined security boundary and only allow specific computers to communicate with your certification authority. This means clients cannot connect directly to a CA to request and retrieve certificates. Or you may want to set up a CA that supplies certificates to Internet-based subjects. You may want to place a layer of insulation between these clients and the CA. Web enrollment allows these scenarios to be implemented easily.

Web enrollment is the component of a Microsoft CA that allows clients to submit certificate issuance and retrieval requests based on Web pages. It is installed when you install a Microsoft CA. It can also be installed on a non-CA computer to provide a Web front-end to your CA infrastructure on the back-end. In this configuration, clients make requests with the Web pages and only the computer that has Web enrollment installed needs to communicate with your CA.

For more information on installing Web enrollment, see Set up certification authority Web enrollment support.

Public and private key pair: cryptographic service providers, key lengths, hash algorithms

If you select the **Use custom settings to generate the key pair and CA certificate** check box when you choose the type of CA to install, you can select the cryptographic service provider (CSP) to use. The CSP generates the public key and private key pair and performs cryptographic operations on behalf of the CA.

You can set the key length for the public key cryptographic keys that the CA uses to sign certificates. In general, the longer the key length, the more secure it is. Note that a longer key will take more time to generate during setup and longer to transmit on the network when building certificate trust lists.

You can also choose the message hash algorithm used by the CA, as well as specifying the use of existing cryptographic keys instead of generating new ones.

Note

- A CA may be installed using a smart card CSP. Once a CA is installed using a smart card CSP, the smart card must be available and inserted with the personal identification number (PIN) in order for the CA to start and perform its operations. This is a higher security option for CA operations than other options, but a disadvantage of using smart card CSPs is that CA performance may be slower and key length may be smaller.

Certification authority identifying information

The following are some guidelines for completing the CA identifying information in Certificate Services setup:

Field	Description
CA Name	The name you want to give to the CA. You can enter a string using any Unicode character. The name of the CA will also be the common name of the CA's distinguished name in Active Directory. When special characters exist in the CA name, a sanitized CA name is used for operations that are unable to use the unmodified CA name. A CA's sanitized name is the name of the CA with all special characters encoded in a form that will allow them to be used for file names, CryptoAPI key container names, and Active Directory object names. Special characters are those characters that cannot be used in one or more of these names; the list includes all characters which are not ASCII characters and many ASCII punctuation characters. Further, Active Directory object names are limited to 64 characters by the Lightweight Directory Access Protocol (LDAP) standard. To accommodate this limit, Active Directory object names are constructed by truncating the sanitized name and appending a hash computed over the truncated part of the sanitized name. The distinguished name suffix field is automatically filled with the distinguished name of the Active Directory domain. If you edit this distinguished name, you must conform with the LDAP standard. Type **certutil.exe** at a command prompt without arguments to see the sanitized name for all of the published CAs. Type **certutil.exe -v -ds** to see all of the CA-related Active Directory names. The first column is a truncated CA name with the hash appended (the actual Active Directory object's container name, with special characters reverted back to their original form). A second column is displayed only if the truncated, sanitized name does not match the truncated CA name, and it is the actual Active Directory object's container name.
Distinguished name suffix	The X.500 distinguished name suffix that will be appended to the CA name.
Validity period	The length of time the CA's certificate should be valid. The CA will obtain a certificate that is valid for this length of time and use that certificate's private key for signing issued certificates and certificate revocation lists.

Database and configuration storage

Certificate Services uses local storage for its database, configuration data, backup data and logging data. You can specify locations for the database and log during CA setup. By default, the certificates issued by a CA are stored in:

Systemroot\system32\certlog

For best performance, the database and log files should be kept on separate physical disk drives, preferably on separate disk controllers. This will maximize disk throughput and allow the CA to perform better.

You also have the option of specifying a shared folder when setting up a CA. The shared folder acts as a location where computer users can find information about certification authorities. This option is only useful if you are installing a stand-alone CA and do not use Active Directory.

If the host computer for Certificate Services is a member of a domain, information about the CA is automatically published to the Active Directory. However, the Active Directory does not act as the database for Certificate Services. This function remains with the local computer.

Note

- It is considered unsafe to install Certificate Services on an FAT file system because FAT does not support domain-based security. Always install Certificate Services on a NTFS file system to take advantage of support for a variety of features including Active Directory, which is needed for domains, user accounts, and other important security features.

(Optional) Creating an issuer policy statement for the CA

When you set up a CA, you can add a CA policy statement to the CA certificate that is created during setup or a CA certificate renewal, in the form of text or a pointer to a Web site. The CA policy statement gives legal and other pertinent information about the CA and its issuing policies, limitations of liability, and so on. An end user will see this CA policy statement when they view the CA certificate and click **Issuer Statement**.

Here are the steps you need to follow to attach a policy statement to a CA's certificate:

- The policy statement file must be installed *before* you set up Certificate Services. This file, named CAPolicy.inf, must be placed in the systemroot directory. URLs in CAPolicy.inf should use the replaceable parameter syntax that also appears in the **Extensions** tab. If you don't use the replaceable parameter syntax, then CDP and AIA extensions in renewed CA root certificates may point to the same location that was specified in the previous CA root certificate. To see the syntax for replaceable parameters, see Specify certificate revocation list distribution points in issued certificates.

- The first two lines of CAPolicy.inf must be:

  [Version]

  Signature="$Windows NT$"

- The next two lines of the file list the name of the policies for this CA. Multiple policies can be listed in **Policies=** if they are separated by commas. The name *LegalPolicy* is used here as an example, but the name can be whatever the CA administrator chooses when creating CAPolicy.inf:

  [CAPolicy]

  Policies=*LegalPolicy*

- For each policy, you need to provide a user-defined object identifier (also known as an OID) and either the text you want displayed as the policy statement or a URL pointer to the policy statement. (The URL can be in the form of an HTTP, FTP, file, or LDAP address.) Continuing on with the example, if you are going to have text in the policy statement, the next three lines of CAPolicy.inf will read:

  [*LegalPolicy***]**

OID=*1.1.1.1.1.1.1.1.1*

Notice="*Legal policy statement text***"**

If you are going use a URL to host the CA policy statement, the next three lines would instead read:

[*LegalPolicy***]**

OID=*1.1.1.1.1.1.1.1.1*

URL="http://*CompanyWebSite/CAPolicy/default.asp***"**

(Please note that the "OID=" entry used in this example is arbitrary and is shown only for illustrative purposes.)

In addition:

- Multiple URLs are supported
- Multiple notices are supported
- Mixed notices and URLs are supported
- URLs with spaces or text with spaces must be surrounded by quotes

An example of multiple notices and URLs in a policy section might be:

[*LegalPolicy***]**

OID=*1.1.1.1.1.1.1.1.1*

URL = "http://*http.microsoft.com/somewhere/default.asp* **"**

URL = "ftp://*ftp.microsoft.com/somewhere else/default.asp***"**

Notice = "*Legal policy statement text.***"**

- You can specify CRL Distribution Points (CDPs) in CAPolicy.inf. Note that any CDP in CAPolicy.inf will take precedence for certificate verifiers over the CDP's specified in the CA policy module. If you want to specify the CDP using CAPolicy.inf, a sample of the syntax is:

  [CRLDistributionPoint]

  URL="http://*CompanyWebSite/Public/myCA.crl***"**

  Some additional notes about the CRL section:

- Multiple URLs are supported.
- HTTP, file, FTP, and LDAP URLs are supported.

  - This section can be used during CA setup or CA certificate renewal.
  - This section is only used if you are setting up a root CA or renewing a root CA certificate. Subordinate CA CRL distribution point extensions are determined by the CA which issued the subordinate CA's certificate.
  - URLs with spaces must be surrounded by quotes.
  - If no URLs are specified (in other words, if the **[CRLDistributionPoint]** section is empty), the default CRL Distribution Point extension will be suppressed.

- You can specify the authority information access points in CAPolicy.inf. The syntax is:

  [AuthorityInformationAccess]

  URL="http://*CompanyWebSite/Public/myCA.crl***"**

  Some additional notes about the authority information access section:

  - Multiple URLs are supported.
  - HTTP, file, FTP, and LDAP URLs are supported.
  - This section can be used during CA setup or CA certificate renewal.
  - This section is only used if you are setting up a root CA or renewing a root CA certificate. Subordinate CA authority information access extensions are determined by the CA which issued the subordinate CA's certificate.
  - If no URLs are specified (in other words, if the **[AuthorityInformationAccess]** section is empty), the default Authority Information Access extension will be suppressed.

- Another optional section of CApolicy.inf is **[EnhancedKeyUsage]**, which is used to specify Enhanced Key Usage extension object identifiers (also known as OIDs).

  - Multiple object identifiers are supported.
  - This section can be used during CA setup or CA certificate renewal.
  - This section is only used if you are setting up a root CA or renewing a root CA certificate. The Enhanced Key Usage extension for a subordinate CA is determined by the CA which issued the subordinate CA's certificate.

  An example of this section is:

  [EnhancedKeyUsage]

  OID=*1.2.3.4.5*

  OID=*1.2.3.4.6*

- Another optional section of CApolicy.inf is **[certsrv_server]**, which is used to specify renewal key length, the renewal validity period, and the certificate revocation list (CRL) validity period for a CA that is being renewed or installed.

  An example would be:

  [certsrv_server]

  RenewalKeyLength=*2048*

  RenewalValidityPeriod=*Years*

  RenewalValidityPeriodUnits=*5*

  CRLPeriod=*Days*

  CRLPeriodUnits=*2*

  CRLDeltaPeriod=*Hours*

  CRLDeltaPeriodUnits=*4*

- **RenewalKeyLength** sets key size for renewal only. This is only used when CA renewal is generating new keys.

- **RenewalValidityPeriod** and **RenewalValidityPeriodUnits** establish the lifetime of new root CA certificate when renewing the old CA certificate.

- **CRLPeriod** and **CRLPeriodUnits** establish the validity period for the full CRL, while **CRLDeltaPeriod** and **CRLDeltaPeriodUnits** establish the validity period for the delta CRL.

- **RenewalKeyLength**, **RenewalValidityPeriod** and **RenewalValidityPeriodUnits** are only processed when the CA is being renewed.

- **LoadDefaultTemplates** is not processed when the CA is a subordinate CA.

Revocation list

From Wikipedia, the free encyclopedia
Jump to: navigation, search

In the operation of some cryptosystems, usually public key infrastructures (PKIs), a **certificate revocation list (CRL)** is a list of certificates (or more specifically, a list of serial numbers for certificates) that have been revoked, and therefore should not be relied upon.

Contents

[hide]

[edit] Revocation States

There are two different states of revocation defined in RFC 3280:

- Revoked: A certificate is irreversibly revoked if, for example, it is discovered that the certificate authority (CA) had improperly issued a certificate, or if a private-key is thought to have been compromised. Certificates may also be revoked for failure of the identified entity to adhere to policy requirements such as publication of false documents, mis-representation of software behavior, or violation of any other policy specified by the CA operator or its customer. The most common reason for revocation is the user no longer being in sole possession of the private key (*e.g.*, the token containing the private key has been lost or stolen).
- Hold: This reversible status can be used to note the temporary invalidity of the certificate (*e.g.*, if the user is unsure if the private key has been lost). If, in this example, the private key was found and nobody had access to it, the status could be reinstated, and the certificate is valid again, thus removing the certificate from future CRLs.

[edit] Reasons for Revocation

Reasons to revoke a certificate according to RFC 5280 p69 are:

- unspecified (0)
- keyCompromise (1)
- cACompromise (2)
- affiliationChanged (3)
- superseded (4)
- cessationOfOperation (5)
- certificateHold (6)
- (value 7 is not used)
- removeFromCRL (8)
- privilegeWithdrawn (9)
- aACompromise (10)

[edit] Publishing Revocation Lists

A CRL is generated and published periodically, often at a defined interval. A CRL can also be published immediately after a certificate has been revoked. The CRL is always issued by the CA which issues the corresponding certificates. All CRLs have a lifetime during which they are valid; this timeframe is often 24 hours or less. During a CRL's validity period, it may be consulted by a PKI-enabled application to verify a certificate prior to use.

To prevent spoofing or denial-of-service attacks, CRLs usually carry a digital signature associated with the CA by which they are published. To validate a specific CRL prior to relying on it, the certificate of its corresponding CA is needed, which can usually be found in a public directory.

The certificates for which a CRL should be maintained are often X.509/public key certificates, as this format is commonly used by PKI schemes.

[edit] Revocation vs. Expiration

Certificate expiration dates are not a substitute for a CRL. While all expired certificates are considered invalid, not all unexpired certificates are necessarily valid. CRLs or other certificate validation techniques are a necessary part of any properly operated PKI, as mistakes in certificate vetting and key management are expected to occur in real world operations.

In a noteworthy example, a certificate for Microsoft was mistakenly issued to an unknown individual, who had successfully posed as Microsoft to the CA contracted to maintain the ActiveX 'publisher certificate' system (VeriSign)[1]. Microsoft saw the need to patch their cryptography subsystem so it would check the status of certificates before trusting them. As a short-term fix, a patch was issued for the relevant Microsoft software (most importantly Windows) specifically listing the two certificates in question as "revoked".[2]

[edit] Problems with CRLs

Best practices require that wherever and however certificate status is maintained, it must be checked whenever one wants to rely on a certificate. Failing this, a revoked certificate may be incorrectly accepted as valid. This means that to use a PKI effectively, one must have access to current CRLs. This requirement of on-line validation negates one of the original major advantages of PKI over symmetric cryptography protocols, namely that the certificate is "self-authenticating". Symmetric systems such as Kerberos also depend on the existence of on-line services (a key distribution center in the case of Kerberos).

The existence of a CRL implies the need for someone (or some organization) to enforce policy and revoke certificates deemed counter to operational policy. If a certificate is mistakenly revoked, significant problems can arise. As the certificate authority is tasked with enforcing the operational policy for issuing certificates, they typically are responsible for determining if and when revocation is appropriate by interpreting the operational policy.

The necessity of consulting a CRL (or other certificate status service) prior to accepting a certificate raises a potential denial-of-service attack against the PKI. If acceptance of a certificate fails in the absence of an available valid CRL, then no operations depending upon certificate acceptance can take place. This issue exists for Kerberos systems as well, where failure to retrieve a current authentication token will prevent system access. No comprehensive solutions to these problems are known, though there are multiple workarounds for various aspects, some of which have proven acceptable in practice[citation needed].

An alternative to using CRLs is the certificate validation protocol known as Online Certificate Status Protocol (OCSP). OCSP has the primary benefit of requiring less network bandwidth, enabling real-time and near real-time status checks for high volume or high value operations.

[edit] See also

- Authority revocation list
- Trusted third party
- Web of trust
- Certificate server

[edit] References

1. ^ http://news.cnet.com/2100-1001-254586.html
2. ^ http://www.microsoft.com/technet/security/bulletin/MS01-017.mspx

[edit] External links

- RFC 3280
- RFC 5280

a network of more than 300 machines all running on Windows Server 2003 and clients running Windows 2000 Professional and Windows XP Professional, Microsoft Certification Server in enterprise CA mode is used for issuing and revoking certificates for all users and computers. An employee who left the company last week was using a digital certificate for secure email communication and must have his certificate revoked,since it will no longer be in use. What procedure(s) must be followed to revoke this certificate? All help is greatly appreciated, thanks in advance.

Microsoft, VeriSign, and Certificate Revocation

Written: 20 Apr 2001 -- Last revised: 13 May 2001

Background

In late January 2001, VeriSign erroneously issued two Class 3 code-signing certificates to someone falsely claiming to represent Microsoft. The certificates were issued in Microsoft's name, specifically "Microsoft Corporation". After issuing the certificates, a routine VeriSign audit uncovered the error in mid-March, about 6 weeks later. VeriSign then did three things:

1. It notified Microsoft of the error.
2. It posted a public notice.
3. It revoked the certificates in its normal Certificate Revocation List (CRL)

Microsoft's immediate response was to announce the problem and describe what users could do to actively recognize and avoid trusting the fraudulent certificates. Microsoft's ultimate response was to issue a patch, described in Microsoft Security Bulletin MS01-017.

It's important to understand that the two certificates at the center of this incident are **real VeriSign certificates**, not fakes masquerading as real. The fundamental error is that VeriSign issued those certificates to somebody other than Microsoft, so somebody can then create and sign software that appears to come from Microsoft.

It's also important to understand that VeriSign's normal procedure for revoking **ANY** of its certificates is to publish the certificate serial-numbers in its CRL. VeriSign maintains many CRLs, one for each class and category of its certificates. You can see a list of all the available CRLs simply by visiting VeriSign's CRL home.

Stirring Up Hornets

Bruce Schneier, a well-known and widely respected security expert, then wrote about the episode in the April issue of his monthly newsletter Crypto-Gram, in an article entitled **Fake Microsoft Certificates**. In that article, Bruce made several statements about the nature of Microsoft's security infrastructure. One statement was:
[...] there is no way to revoke the certificates (Windows has no CRL features), [...]

By coincidence, a letter I wrote to Crypto-Gram about a previous article also appeared in the same issue, and I specifically cited the bogus certificate episode in one of my examples. I also made a statement about Microsoft's security infrastructure, specifically:

[...] no Microsoft software is capable of automatically obtaining the Certificate Revocation List (CRL) listing those two bogus certificates, because there's no revocation infrastructure.

Microsoft did not agree with either one of these characterizations. Two things then happened:

1. Late on 14 April 2001, the program manager of .NET security, David B. Cross, sent me an email containing his argument against my characterization, ending with:
   "*I believe your statement was made in error and should be corrected.*"
   I did not accept his argument, and we engaged in a few responses and counter-responses. I will not present his specific argument here, because it essentially mirrored the argument that would appear later in a public Microsoft response.
2. About 16 April 2001, Microsoft published its Response to Inaccurate Crypto-Gram Article on VeriSign Certificates. Among other things, it specifically claimed that there was in fact a way to revoke certificates using a CRL in Windows:

   There is a way to revoke the certificates -- via the Certificate Revocation List (CRL) mechanism defined in the relevant industry standard, RFC 2459. And Windows does indeed have CRL features -- CRL support has been available for the Windows NT family since 1998, and for the Windows 9x family since early 1999.

What follows is essentially a rebuttal to Microsoft's latest statements, and explains why, in any practical sense, Windows does not have a revocation infrastructure capable of revoking the two VeriSign certificates in question.

At Issue

The issue is not whether Microsoft provided a patch. It plainly did.

The issue is not whether the patch is effective in blocking the two certificates or or not. It presumably is.

The specific point at issue here is **whether Windows, via its CryptoAPI, does or does not have a viable revocation infrastructure.** We can be even more specific and confine ourselves to the viability of the revocation infrastructure just for **VeriSign's Class 3 code-signing certificates**, since those are the kind of certificates at the center of this incident. They are also the most likely kind of certificates that developers will buy and sign their code with.

This may seem to be a minor quibble, but revocation is a crucial part of any certificate-based security infrastructure. That's because **certificates can be revoked for any number of reasons prior to their expiration date**.

Issuing and Revoking Certificates

Any organization that issues certificates is called a **Certificate Authority** (CA). As a rule, every certificate issued by a CA contains an expiration date (expiry), placed directly into the certificate by the CA when the certificate is made. Each certificate also contains its own class and category, e.g. a class 3 code-signing certificate, so anyone who validates certificates can tell exactly what privilege or capability should be granted to the certificate's holder.

Every CA that issues certificates will also need to revoke them. There can be many reasons for this. One reason is that the private keys corresponding to the certificate have been compromised or lost. Then the certificate holder contacts the CA and asks that the certificate be revoked. The CA authenticates this request before actually listing the certificate on its CRL. A CA itself may also decide to revoke a certificate. For example, if the certificate holder violates the issuing agreement, or if the certificate was issued in error, such as happened with the bogus Microsoft certificates. Whatever the reason, the certificate's serial-number appears on the CRL of the issuing CA, and anyone who receives the certificate should decline to trust it.

Each CA decides how long its certificates will last, and customers know the length of this period (typically one year) before they buy the certificate. The expiry cannot be modified without making the certificate invalid. Since the expiry is right there in the certificate, anyone reading the certificate to determine its trustworthiness can determine the expiry for themselves. They don't need to be told by some other source when any given certificate will expire.

Things are not so easy when a certificate must be revoked before its normal expiry. Then, anyone reading the certificate needs additional data: the Certificate Revocation List (CRL) corresponding to the certificate. CA's typically provide a CRL at some well-known and publically accessible location on the Internet, i.e. at some URL. CA's also update their CRL's when new revocations occur. Each CRL is digitally signed and dated by the issuing CA, so everyone having the CRL knows it was really generated by the CA rather than some random interloper.

CA's typically issue CRLs at some periodic rate, whether there are changes or not. This can be daily or weekly in typical cases. It may be more or less frequently, according to the CA's policies. The reason for revocation, hence the rationale for the update frequency may also vary. The key phrase used to describe the timeliness of CRLs is **suitably-recent** and can be found in RFC 2459:

```
3.3  Revocation

   When a certificate is issued, it is expected to be in use for its
   entire validity period.  However, various circumstances may cause a
   certificate to become invalid prior to the expiration of the validity
   period. Such circumstances include change of name, change of
   association between subject and CA (e.g., an employee terminates
   employment with an organization), and compromise or suspected
   compromise of the corresponding private key.  Under such
   circumstances, the CA needs to revoke the certificate.
     [...]  When a
   certificate-using system uses a certificate (e.g., for verifying a
   remote user's digital signature), that system not only checks the
   certificate signature and validity but also acquires a suitably-
   recent CRL and checks that the certificate serial number is not on
   that CRL.  The meaning of "suitably-recent" may vary with local
   policy, but it usually means the most recently-issued CRL.
   [...]
```

In short, CAs decide how often to issue CRLs, and users of those certificates ultimately decide how often to obtain a suitably-recent CRL. This does not remove the need for the CRL. That is, "never" is not an allowable value for "suitably-recent".

All certificate users, i.e. anyone who examines certificates to determine its trustworthiness, need all the following pieces of information in order to determine the complete trustworthiness of any certificate:

1. The certificate in question, i.e. the one whose trust is being evaluated.
2. The already trusted **root certificate** of the CA that issued the certificate in question.
3. A **suitably-recent CRL** issued by the CA covering the class and category of the certificate in question.

The trusted root CA certificate lets you determine the trustworthiness (authenticity) of both the certificate in question and the CRL. The certificate in question contains its own expiry, class, and category, so you can determine those yourself after verifying the authenticity of the certificate itself. If you don't have a CRL, however, then you have incomplete information. You can tell that the certificate in question was issued by the CA, and was valid at some point and for some time after it was issued, but you can't tell anything about whether that certificate is still valid. The CA could have revoked the certificate and you wouldn't know. You'd end up trusting a certificate that you shouldn't have.

So where do you obtain all these three elements? The trusted root CA certificate is something you already have. It has to be, because you have to trust it already. The certificate in question is presented to you each time. But where do CRLs come from? Ultimately from a CA, but how?

Revocation Infrastructure

CRLs are obtained using a revocation infrastructure. This just means that for every CA you trust, and whose root CA certificate you use, you have a corresponding URL to download a CRL, or some other means to obtain one. On its face, this seems simple enough, and it is. The difficulties come in how a particular CA tells you where to obtain its CRLs, and whether your computer does anything about it. This is also the basis for Bruce's and my assertions that Windows has no revocation infrastructure.

When a CA tells you where to obtain CRLs, it must do so in a trustworthy manner. That is, it must be effectively impossible for any random interloper to provide a fake, outdated, or otherwise bogus CRL. If you used a fake or outdated CRL, someone could convince you a certificate was valid when it really wasn't.

Some things that CA's do to ensure authentic CRLs are:

1. Every CRL is digitally signed by the CA using its root certificate. This ensures that CRLs can't be faked.
2. Every CRL is dated by the CA. This ensures that every CRL eventually becomes stale or expires and you obtain the latest one.

3. Every CRL has a higher sequence number than the one issued before it. This ensures that no one can slip you a CRL earlier than the one you have now, even one that has not yet expired.

All these checks are part of the CRL itself. But this still doesn't tell you where or how to obtain a CRL in the first place. As it turns out, there can be many different ways, none of which is actually **mandated** by any standard:

- **The Manual Method** -- You could make users manually download a CRL from a certain URL, install it in a specific place on the local machine, and then use the local file. This has the same disadvantage as every manually maintained security update or fix: users don't do it.
- **The Well-Known-URL Method** -- Since you already have a trusted root certificate for every CA, you can also have a specific known and trusted URL associated with every CA. Every CA would publish this URL in a well-known location, so everyone would know it. When you need a CRL, you just look up a URL in your trusted list of associations, retrieve the CRL, and you're ready. ("Well-known" primarily means well-known to software developers, since they create the default trusted associations, though users can update or edit them.)
- **The Root-Certificate-URL Method** -- Every CA could write the URL of its current CRL in its trusted root certificate. Then when you need the CRL you just read its URL from the trusted root certificate, retrieve the CRL, and you're ready. (This is actually just a simple variation on the next method.)
- **The RFC 2459 Method** -- Rather than writing a URL only in its root certificate, every CA could write the URL of a CRL into every certificate it issues. The advantage of this over the Root-Certificate-URL method is that it allows a CA to add new URLs over time. RFC 2459 is the document where this method is described.

There are certainly other methods I'm omitting, including real-time on-line certificate status checks, but my purpose is to list some workable examples, not to create an exhaustive list.

However CRLs are distributed or obtained, there are several critical points about them, which may seem obvious in their simplicity:

1. You must know where to obtain a suitably-recent CRL.
2. You must actually obtain that CRL.
3. You must authenticate the contents of the CRL.
4. You must actually use the CRL.

Omit any of these, and it doesn't matter what else happens, you won't have a working revocation infrastructure. The whole point of a revocation infrastructure is to ensure that valid up-to-date CRLs are available when and where they are needed.

Publishing CRLs

Now that you know what a revocation infrastructure is, and why it's important, let's return to the VeriSign/Microsoft case. We already know that Windows accepts VeriSign certificates.

Microsoft made that decision for all its users by including VeriSign's trusted root certificate in its product. This means Windows should also be obtaining VeriSign's CRLs and checking VeriSign's certificates against those CRLs. If it doesn't, then Windows users will unwittingly accept and trust certificates already revoked by VeriSign. Remember that revocations can occur for many reasons, including lost or compromised keys, violations of VeriSign policy, etc. and not just erroneously issued certificates. Also remember that VeriSign alone revokes VeriSign certificates, not Microsoft. While Microsoft can request that VeriSign revoke one or more of Microsoft's own certificates, Microsoft cannot request that anyone else's certificates be revoked. Well, it can, but VeriSign should not actually revoke any such certificate unless Microsoft presents very good justification to VeriSign that a significant enough infraction has actually occurred.

So exactly which method does VeriSign use to publish its CRLs? And exactly how does Windows obtain and use those VeriSign CRLs?

First, let's examine how VeriSign publishes its CRLs.

Does VeriSign use the RFC 2459 method? According to Microsoft, and we have no reason to doubt them on this point, VeriSign does not. The erroneously issued certificates definitely don't list any CRL locations (also called CRL Distribution Points or CDPs). You can even check this with any program that decodes and dumps X.509v3 certificates, so you don't have to take Microsoft's word for it.

Does VeriSign use the Root-Certificate method? No, there are no CRLs in VeriSign's Class 3 code-signing root certificate, either.

Are VeriSign's CRLs available from some well-known URL? Yes. While ordinary users may not consider this a "well-known" URL, developers of security products that accept VeriSign certificates should.

Are VeriSign's CRLs manually retrievable? Yes, you can download any of them. with any browser or other HTTP client.

Clearly, VeriSign is leaving it to the actual crypto-library developers to either retrieve a CRL from its well-known URL, or to somehow give users a means to manually install a VeriSign CRL after it's been manually downloaded. The crypto-library developers for Windows is Microsoft, who provides the CryptoAPI that all its code-signing security is built on. Yes, third parties can provide additional modules and plug-ins, but the entire basic capability is defined and provided by Microsoft.

RFC 2459

Now let's look more closely at RFC 2459. Is VeriSign violating a standard by not providing a CRL location in every issued certificate, as RFC 2459 presumably states? Given the Microsoft rhetoric, the rather surprising answer is "No". In fact, RFC 2459 says that giving a CRL location within a certificate is optional. The exact relevant quote is this:

```
4.2.1.14  CRL Distribution Points

   The CRL distribution points extension identifies how CRL information
   is obtained.  The extension SHOULD be non-critical, but this profile
   recommends support for this extension by CAs and applications.
     [...]
```

The meaning of "non-critical" appears earlier in RFC 2459:

```
4.2  Standard Certificate Extensions

     [...]  Each extension in a
   certificate may be designated as critical or non-critical.  A
   certificate using system MUST reject the certificate if it encounters
   a critical extension it does not recognize; however, a non-critical
   extension may be ignored if it is not recognized.
     [...]
   Conforming CAs MUST support key identifiers (see sec. 4.2.1.1 and
   4.2.1.2), basic constraints (see sec. 4.2.1.10), key usage (see sec.
   4.2.1.3), and certificate policies (see sec. 4.2.1.5) extensions.
     [...]
   Support for the remaining extensions is OPTIONAL.
```

In other words, the words "critical" and "non-critical" refer to how a certificate should be handled when a field within the certificate is actually marked "critical" **in a specific certificate**, not to whether a particular extension defined in RFC 2459 is required or optional. Criticality is a way to ensure that new extensions are treated with proper consequences when evaluated by old software that may not recognize the meaning or uses of the extension. Marking the new extension as critical ensures that old software will reject such certificates, thereby preserving the overall security and integrity of what the certificates are protecting.

The precise meaning of the capitalized words SHOULD, MUST, and OPTIONAL is given in RFC 2119:

```
   MUST   This word, or the terms "REQUIRED" or "SHALL", mean that the
   definition is an absolute requirement of the specification.

   SHOULD   This word, or the adjective "RECOMMENDED", mean that there
   may exist valid reasons in particular circumstances to ignore a
   particular item, but the full implications must be understood and
   carefully weighed before choosing a different course.

   MAY   This word, or the adjective "OPTIONAL", mean that an item is
   truly optional.  One vendor may choose to include the item because a
   particular marketplace requires it or because the vendor feels that
   it enhances the product while another vendor may omit the same item.
   An implementation which does not include a particular option MUST be
   prepared to interoperate with another implementation which does
   include the option, though perhaps with reduced functionality. In the
   same vein an implementation which does include a particular option
   MUST be prepared to interoperate with another implementation which
   does not include the option (except, of course, for the feature the
   option provides.)
```

One may question whether VeriSign has "understood and carefully weighed" the full implications, but that's a completely different question than whether they followed RFC 2459 or not. The plain and simple fact is that VeriSign did not provide the CRL distribution point extension of RFC 2459, and it was within standardized practice to omit it. If Microsoft objected to this omission, it was entirely free to omit VeriSign's Class 3 code-signing root certificate from its product, and prevent all such certificates issued by VeriSign from being accepted. No one forced Microsoft to use VeriSign certificates lacking the CRL extension. Microsoft alone chose to do so.

Obtaining VeriSign CRLs

We've seen how VeriSign publishes its CRLs, and examined whether its certificates and CRLs are within RFC 2459's requirements. Now let's revisit the question of exactly how Windows obtains and uses VeriSign CRLs. In particular, does Microsoft's CryptoAPI support the Well-Known-URL method, or does it rely on the user to provide the Manual method? The rather astonishing conclusion one must arrive at is "Neither".

To be specific, **Microsoft's CryptoAPI, as shipped by Microsoft, only handles CRLs when they are listed in certificates that have the CRL Distribution Point extension of RFC 2459.**

At least this is the only conclusion one can draw from Microsoft's public responses to the incident. You can see this for yourself by looking at the Response to Inaccurate Crypto-Gram Article on VeriSign Certificates, searching for the string "2459", and reading the nearby "explanation":

There is a way to revoke the certificates -- via the Certificate Revocation List (CRL) mechanism defined in the relevant industry standard, RFC 2459.

In fact, Microsoft makes a big deal about RFC 2459 in its response, and how the certificates don't list a CRL location. It even goes so far as to include a screen shot showing that the certificates don't have the RFC 2459 CRL Distribution Point extension in them.

But no one ever said they did.

No one is saying that now.

This is all a red herring. Microsoft is simply evading the real question of how it could design and ship a revocation infrastructure relying entirely on a feature it must have known didn't exist in the VeriSign certificates it was accepting.

A similar "explanation" can be seen in Microsoft Security Bulletin MS01-017, after the first occurrence of the string "CRL":

However, because VeriSign's code-signing certificates do not specify a CRL Distribution Point (CDP), it is not possible for any browser's CRL-checking mechanism to locate and use the VeriSign CRL. Microsoft has developed an update that rectifies this problem. The update package includes a CRL containing the two certificates, and an installable revocation handler that consults the CRL on the local machine, rather than attempting to use the CDP mechanism.

Also, the bulletin ends with the assertion:
Without this data, no system -- Microsoft or otherwise -- could obtain and check the CRL.
Again, Microsoft is basically saying, "If only VeriSign certificates had conformed to RFC 2459 and included a CRL location (or CDP), everything would have worked fine". It's also saying that no other browser or security infrastructure can obtain or use a VeriSign CRL either.

On it's face, though, this new assertion is absurd. Any system -- even Microsoft's -- could have an internal Well-Known URL for VeriSign, from which it automatically obtains suitably-recent CRLs for any VeriSign certificates it cares to accept. As an unautomated alternative, any system -- even Microsoft's -- could have a simple Manual procedure for obtaining and using CRLs. Which systems actually do this is a different question entirely. Indeed, one system we certainly know cannot and does not: Microsoft's.

But maybe a simple Manual process **WILL** work with MSIE. Here's what one user wrote to me about the process:

1. Go to the VeriSign CRL site and click on the link named **Class3SoftwarePublishers.crl** to download the CRL.
2. Look in the MSIE Help menu to find instructions about what to do with it:
   a. Go to Tools | Internet Options, then click on the Contents Tab.
   b. Click the Certificates button.
   c. Click Import, then use the Certificate Manager Import Wizard to import the CRL.
3. Follow the confirmation instructions in Microsoft Security Bulletin MS01-017:

   In IE, select Tools, then Internet Options. Select the Advanced tab, then scroll to the section titled Security and verify that "Check for publisher's certificate revocation" has been selected.

4. Verify that the CRL is in use by downloading an OCX and following the instructions on Microsoft's update confirmation page.
5. [But it's not working because] I followed the procedure listed, and what it says under Digital Signature Information is "This signature is okay" rather than anything about revocation.

   Now it's possible that I've done something wrong, but if I have, it's something that a very computer-literate person (though admittedly not a techie) managed to miss. Either way, something is very wrong.

So maybe Microsoft has a working Manual procedure and maybe they don't -- it seems to do something but we can't confirm it. And however we might characterize the process, it can hardly be called simple or obvious (unlike, say, downloading JPEG images or Windows media files). Complexity and obscurity are the enemies of effective security. If users can't confirm that the Manual process has worked and is effective, why should they trust it? And why should they consider it to be a working revocation infrastructure?

Whether Microsoft's assertion that "no system ... could obtain and check the CRL" is valid in any practical sense or not, it is a hollow one. It is at best wishful thinking, since it ignores the concrete realities of VeriSign's actual certificates and CRLs. The really interesting thing is that VeriSign has been publishing CRLs under its current mechanisms **FOR YEARS**, indeed, for years before RFC 2459 was ever written. Either it was wasting time and money doing so, since those CRLs were never used by anyone, or someone somewhere managed to figure out a way to get the VeriSign CRLs and use them properly. That Microsoft was unable to do so is the real issue here.

We've already seen that **the CRL extension in RFC 2459 is OPTIONAL**, even though recommended. We've also seen that VeriSign does not issue Class 3 code-signing certificates containing CRL locations, and apparently never has. Whether it does in the future is irrelevant. There are countless Class 3 code-signing certificates already issued and deployed in all kinds of software. All that software will stop working unless Microsoft provides something like an automatic Well-Known-URL method as part of the way its revocation infrastructure handles VeriSign CRLs.

This leaves a really big unanswered question:

Why would Microsoft base its entire CRL and revocation infrastructure on an optional feature that was absent from certificates issued by one of its principal certificate providers?

At the very least, there are some gravely flawed design assumptions here, or there is a major breakdown in communication and requirements. Microsoft alone knows the real answer, and Microsoft alone should be held accountable.

Amazingly, this condition (unobtainable and/or unusable VeriSign CRLs) must have persisted for some time. That is, the revocation problem itself has existed since CryptoAPI first shipped. Windows has **NEVER** been able to obtain and use a VeriSign CRL, at least for any practical meaning of the words "obtain" or "use". If it had, then Microsoft would not have had to issue **ANY** patch or update. All it would have had to do was let the normal revocation infrastructure work in the normal way, and all Windows users would eventually have received VeriSign's normally published CRL using the normal revocation infrastructure. Even if Microsoft had a clearly working Manual method for obtaining CRLs, all it would have had to do was tell its customers to obtain a new VeriSign CRL. After all, that's the whole point of having an infrastructure in the first place -- to avoid continually reinventing wheels or blazing new trails through wilderness each and every time. Whatever else would have been necessary, it is certain that no uniquely issued and manually installed Microsoft patch or "update" would have been needed. The infrastructure would have worked exactly the way it was intended to work, whether that was automatically or manually.

Security is a chain. Break one link and the whole chain fails. The critical broken link in this incident was Microsoft's sole reliance on a feature that simply didn't appear in the certificates it was accepting. Microsoft says this still means it has a working revocation infrastructure. But if their principal provider of certificates, VeriSign, doesn't include the optional feature pivotal to their revocation infrastructure, then they are left with an utterly vacuous meaning for "working", much less for "revocation infrastructure".

Infrastructure Fixes

So Microsoft has issued a patch, er, I mean "update", and now Windows programs can obtain and use future VeriSign CRLs automatically, and everything will be fine from here on out.

Or will it? Remember that CRLs are periodically reissued with new entries added and outdated ones removed. It is necessary to have a **suitably-recent CRL**, not just any old CRL from some past period. So how does Windows obtain and install one of the normally issued and suitably-recent VeriSign CRLs?

According to the Microsoft Security Bulletin MS01-017, the patch installs a "private CRL" **containing only three revoked VeriSign certificates**. Two of the revoked certificates are the bogus certificates. The third one is a test certificate that Microsoft specifically revoked so it could test its patch. The "private CRL" does not contain any of the other VeriSign revocations. Why not?

There also appears to be no mechanism in the patch that obtains, installs, or uses any CRL other than the "private CRL" installed by the patch. That "private CRL" is not even a complete VeriSign CRL listing all revoked Class 3 code-signing certificates at the time it was issued. The "private CRL" lists only those certificates that Microsoft deemed to be worthy of revocation. In short, from the moment the patch was issued, it **WASN'T EVEN A SUITABLY-RECENT CRL** since it didn't contain all the revoked certificates, only a Microsoft-selected subset. Why aren't all of VeriSign's other revoked certificates listed? Or has Microsoft determined that accepting any of those other revoked certificates poses no risk to users?

What's going on here? Could it be that Microsoft has simply created an expedient patch that doesn't really fix the larger CRL and revocation infrastructure problem, but merely connects into the CryptoAPI using a CRL plug-in mechanism? Could it be that future VeriSign-issued CRLs will continue to be unknown and unused by Microsoft, just as they always have been?

Or perhaps someday Microsoft will eventually issue a complete fix for VeriSign CRLs, that will always automatically obtain the actual VeriSign-published CRL and use it instead of some "private CRL" installed by a unique manual patch.

We can only hope.

But until that time comes, it seems fairly obvious to conclude that Microsoft does not have a working revocation infrastructure. At least not one that works for one of its principal suppliers of code-signing certificates, VeriSign.

On Certificate Expiration

Microsoft's Response to Inaccurate Crypto-Gram Article on VeriSign Certificates contains another interesting "clarification", this time regarding how long the bogus certificates will last. Bruce Schneier wrote:

Some news reports claim that the certificates will expire in a year, others claim that they will be good forever.

Microsoft's response to this simple statement is:

Fact: It's no mystery when the certificates will expire. The FAQ section of the bulletin provides screen shots of the certificates, and it can plainly be seen that one expires on 30 January 2002, and the other expires on 31 January 2002. This data is corroborated by information on the VeriSign web site.

There are two points to rebut here:

1. No one questioned what the expiration dates of the certificates are;
2. Schneier was simply saying that "Some news reports ... claim that they will be good forever". He was not saying whether these reports were correct or not.

What Schneier should have done is explain why some news reports were saying that the certificates would have been "good forever". So let's examine this question of expired certificates a little more closely.

The short answer that Schneier should have given is that some browsers are perfectly willing to accept expired certificates. In fact, it is not unknown for browsers to accept expired SSL certificates, perhaps with a warning dialog to the user, perhaps not. It is also common for older versions of browsers to actually be relying on expired root certificates. For example, some CA's had root certificates that expired on 31 Dec 1999. Those CA's had also issued countless SSL certificates that also expired on 31 Dec 1999. But on 01 Jan 2000, the world didn't come to a halt. All that happened is that some people visiting SSL sites were asked whether to continue trusting the expired certificates or not. Most people did. Were they complete fools to do so? Were they falling into a security hole? Were they exposing themselves to needless risk? Not necessarily.

Quoting from Schneier's book "Secrets and Lies" (Wiley, 2000), where he discusses the expiration of real-world credentials like passports, credit cards, and driver's licenses:

Expirations provide a safety net. A bad credential can be out there for only so long, because it will expire eventually. *(p. 229)*

But as noted above, and also described at more length in "Secrets and Lies", the credentials we use in the digital world (i.e. certificates) are, in practice, used in ways very different from the ways we use real-world credentials. For one thing, expiration dates on real-world credentials are usually enforced much more strictly. That's not saying that laxity in enforcing digital certificate expiration dates is a good thing, just that it happens regularly and the digital world hasn't come crashing down as a result.

In any case, certificate expirations only provide a safety net. The situation for signed code is not substantially different. If you trusted a code-signing certificate when you first received the code, you really don't need to stop trusting it just because the certificate expires. You decided to trust the certificate at the time you obtained the code, so what changed? Time may have marched on in its petty pace from day to day, but you have exactly the same certificate and the same code, so the trustworthiness of both are no different. What difference does the certificate's expiration make to you? Has your trust of the code's signer or the CA changed in any way between the last day before the certificate expires and the first day after it expires?

As far as trustworthiness is concerned, the critical point in time was **when you first obtained and executed** the signed code, not all the subsequent times you executed it or checked the certificate. If you were justified in deciding to trust the code the first time you ran it, then you were equally justified in trusting it all the other times. Put in the opposite way, if your trust was misplaced the first time you ran the code, then it was equally misplaced all the other times as well. Nothing else about the certificate or the code changes either result, because the fundamental decision has already been made. The basic relationship -- trust the code or not -- has not changed no matter how many other times you evaluate the same certificate and run the same code. Expiration does not change this relationship.

Note that you **CANNOT** apply the same logic to revocation. That is, if a certificate is revoked after you've checked it and chosen to trust it, **that certificate's revocation should make you re-evaluate your trust decision**. In particular, you may want to stop trusting the certificate, depending on the reason for the revocation. For example, if the certificate was revoked because the keys were compromised, then some imposter could have those keys and be able to fool you into trusting some new malicious code. Or maybe the person controlling the certificate's keys has left the company represented by the certificate, and the company revokes the old certificate and has a new certificate issued with new keys. In both cases, something important about the relationship changed, and it could be very risky to continue with the same trust relationship as before the revocation occurred.

Thus, we can see that having a working revocation infrastructure is just as important as having certificate expiration dates. You definitely do not want to be left without a revocation infrastructure when it comes to code-signing. Revocation may even be MORE important than expiration, though I would not want to be without expiration. Doing so would just make certificate management harder.

Certificate expiration is simply the passing of an arbitrary deadline imposed on the certificate by the issuing CA, which may have at least as much to do with a revenue stream for the CA as with any fundamental question of trustworthiness. This is not to denigrate CAs. They take on important risks and perform valuable services. Their business is trust. But it's always worthwhile to understand the relationships as well as the trustworthiness of any business you rely upon. At the very least, it helps you to see latent conflicts of interest.

Open Questions

1. Microsoft says it supports CRLs identified by the CRL Distribution Point field of a certificate defined according to RFC 2459. Windows ships with a significant number of built-in trusted root CA certificates. There may even be updates that include additional built-in trusted root certificates. Of all these built-in trusted root certificates, which CAs provide the CRL Distribution Point feature in their code-signing certificates? Any prudent developer aware of this issue would want to be sure that any code-signing certificate they bought is protected by a CRL that's automatically obtained and used by all Windows machines. So please list the built-in trusted CAs that meet the criterion.
2. For all code-signing certificates that contain an RFC 2459 CRL Distribution Point, is the CRL obtained and checked automatically? Or is this capability enabled and disabled by some user-configurable option? If it's the latter, please explain what I should tell my users to enable, so they know how to enable the revocation infrastructure. Also, please summarize the general effects this will have on bandwidth consumption. That is, are CRLs cached or otherwise managed so they are not downloaded anew with every certificate check?
3. Does Microsoft support subsequent VeriSign CRLs being issued? Say I have a VeriSign code-signing certificate now, but the private keys are compromised somehow. I cancel that certificate, which VeriSign revokes by listing on its CRL. I then buy another VeriSign code-signing certificate which I use thereafter. Will all Windows machines expeditiously learn of and use the new VeriSign CRL listing my revoked certificate? Or is my company effectively shackled to the revoked certificate (and all the mischief that can ensue) until it

actually expires? If it's the latter, then of what practical use is revocation, and how is it possible to call this a revocation infrastructure?

4. Given that the RFC 2459 CRL Distribution Point is optional, is Microsoft planning to update its revocation infrastructure to obtain CRLs by any other means? Just from a practical standpoint, this update would accomodate all those CAs, including VeriSign, who don't have CRL Distribution Points in their certificates.
5. If Microsoft is not planning to update its revocation infrastructure to obtain CRLs by other means, is Microsoft planning to remove the built-in trusted root certificates of CAs who don't conform to the RFC 2459 CRL Distribution Point scheme? True, users could still accept additional root certificates from non-conforming CAs, but shouldn't they at least be told the risks of doing so and allowed to make an informed decision?

To Greg's Home Page
To Greg's Essays & Opinions Page

The Encrypting File System

By Roberta Bragg

On This Page

An Overview of the Encrypting File System

The Encrypting File System (EFS) is a component of the NTFS file system on Windows 2000, Windows XP Professional, and Windows Server 2003. (Windows XP Home doesn't include EFS.) EFS enables transparent encryption and decryption of files by using advanced, standard cryptographic algorithms. Any individual or program that doesn't possess the appropriate cryptographic key cannot read the encrypted data. Encrypted files can be protected even from those who gain physical possession of the computer that the files reside on. Even persons who are authorized to access the computer and its file system cannot view the data. While other defensive strategies should be used, and encryption isn't the correct countermeasure for every threat, encryption is a powerful addition to any defensive strategy. EFS is the built-in file encryption tool for Windows file systems.

However, every defensive weapon, if used incorrectly, carries the potential for harm. EFS must be understood, implemented appropriately, and managed effectively to ensure that your experience, the experience of those to whom you provide support, and the data you wish to protect aren't harmed. This document will

- Provide an overview and pointers to resources on EFS.

- Point to implementation strategies and best practices.
- Name the dangers and counsel mitigation and prevention from harm.

Many online and published resources on EFS exist. The major sources of information are the Microsoft resource kits, product documentation, white papers, and Knowledge Base articles. This paper provides a brief overview of major EFS issues. Wherever possible, it doesn't rework existing documentation; rather, it provides links to the best resources. In short, it maps the list of desired knowledge and instruction to the actual documents where they can be found. In addition, the paper catalogs the key elements of large documents so that you'll be able to find the information you need without having to work your way through hundreds of pages of information each time you have a new question.

The paper discusses the following key EFS knowledge areas:

- What EFS is
- Basic how-tos, such as how to encrypt and decrypt files, recover encrypted files, archive keys, manage certificates, and back up files, and how to disable EFS
- How EFS works and EFS architecture and algorithms
- Key differences between EFS on Windows 2000, Windows XP, and Windows Server 2003
- Misuse and abuse of EFS and how to avoid data loss or exposure
- Remote storage of encrypted files using SMB file shares and WebDAV
- Best practices for SOHO and small businesses
- Enterprise how-tos: how to implement data recovery strategies with PKI and how to implement key recovery with PKI
- Troubleshooting
- Radical EFS: using EFS to encrypt databases and using EFS with other Microsoft products
- Disaster recovery
- Where to download EFS-specific tools

Using EFS requires only a few simple bits of knowledge. However, using EFS without knowledge of best practices and without understanding recovery processes can give you a mistaken sense of security, as your files might not be encrypted when you think they are, or you might enable unauthorized access by having a weak password or having made the password available to others. It might also result in a loss of data, if proper recovery steps aren't taken. Therefore, before using EFS you should read the information links in the section "Misuse and Abuse of EFS and How to Avoid Data Loss or Exposure." The knowledge in this section warns you where lack of proper recovery operations or misunderstanding can cause your data to be unnecessarily exposed. To implement a secure and recoverable EFS policy, you should have a more comprehensive understanding of EFS.

Top Of Page

What EFS Is

You can use EFS to encrypt files stored in the file system of Windows 2000, Windows XP Professional, and Windows Server 2003 computers. EFS isn't designed to protect data while it's transferred from one system to another. EFS uses symmetric (one key is used to encrypt the files) and asymmetric (two keys are used to protect the encryption key) cryptography. An excellent primer on cryptography is available in the Windows 2000 Resource Kit as is an introduction to Certificate Services. Understanding both of these topics will assist you in understanding EFS.

A solid overview of EFS and a comprehensive collection of information on EFS in Windows 2000 are published in the Distributed Systems Guide of the Windows 2000 Server Resource Kit. This information, most of which resides in Chapter 15 of that guide, is published online at http://www.microsoft.com/technet/prodtechnol/windows2000serv/reskit/default.mspx. (On this site's page, use the TOC to go to the Distributed Systems Guide, Distributed Security, Encrypting File System.)

There are differences between EFS in Windows 2000, Windows XP Professional, and Windows Server 2003. The Windows XP Professional Resource Kit explains the differences between Windows 2000 and Windows XP Professionals implementation of EFS, and the document "Encrypting File System in Windows XP and Windows Server 2003" (http://www.microsoft.com/technet/prodtechnol/winxppro/deploy/cryptfs.mspx) details Windows XP and Windows Server 2003 modifications. The section below, "Key Differences between EFS on Windows 2000, Windows XP, and Windows Server 2003," summarizes these differences.

The following are important basic facts about EFS:

- EFS encryption doesn't occur at the application level but rather at the file-system level; therefore, the encryption and decryption process is transparent to the user and to the application. If a folder is marked for encryption, every file created in or moved to the folder will be encrypted. Applications don't have to understand EFS or manage EFS-encrypted files any differently than unencrypted files. If a user attempts to open a file and possesses the key to do so, the file opens without additional effort on the user's part. If the user doesn't possess the key, they receive an "Access denied" error message.
- File encryption uses a symmetric key, which is then itself encrypted with the public key of a public key encryption pair. The related private key must be available in order for the file to be decrypted. This key pair is bound to a user identity and made available to the user who has possession of the user ID and password. If the private key is damaged or missing, even the user that encrypted the file cannot decrypt it. If a recovery agent exists, then the file may be recoverable. If key archival has been implemented, then the key may be recovered, and the file decrypted. If not, the file may be lost. EFS is an excellent file encryption system—there is no "back door."
- File encryption keys can be archived (e.g. exported to a floppy disk) and kept in a safe place to ensure recovery should keys become damaged.
- EFS keys are protected by the user's password. Any user who can obtain the user ID and password can log on as that user and decrypt that user's files. Therefore, a strong password policy as well as strong user education must be a component of each organization's security practices to ensure the protection of EFS-encrypted files.
- EFS-encrypted files don't remain encrypted during transport if saved to or opened from a folder on a remote server. The file is decrypted, traverses the network in plaintext, and, if saved to a folder on the local drive that's marked for encryption, is encrypted locally. EFS-encrypted files can remain encrypted while traversing the network if they're being saved to a Web folder using WebDAV. This method of remote storage isn't available for Windows 2000.
- EFS uses FIPS 140-evaluated Microsoft Cryptographic Service Providers (CSP—components which contain encryption algorithms for Microsoft products).

Top Of Page

Basic How-tos

How to Encrypt and Decrypt Files, Recover Encrypted Files, Archive Keys, Manage Certificates, Back Up Files; and Disable EFS

EFS functionality is straightforward, and you can find step-by-step instructions in many documents online. Links to specific articles for each possible EFS function, as well as some documents which summarize multiple functionality, follow. If the document is a Knowledge Base article, the Knowledge Base number appears in parentheses after the article title.

Encrypting and Decrypting

The process of encrypting and decrypting files is very straightforward, but its important to decide what to encrypt and to note differences in EFS based on the operating system.

- "Encrypting Files in Windows 2000" (222054) explains setting folder encryption. Remember, once a folder is marked for encryption, it isn't necessary to manually mark for encryption the files placed within it.
- "HOW TO: Encrypt a File in Windows XP" (307877) includes error messages and warnings that a user may get when attempting to open files encrypted by another.
- Folders aren't encrypted; however, setting the folder property to "encrypt" does mean that all files placed in the folder will be automatically encrypted "HOW TO: Encrypt a Folder in Windows XP" (308989) tells how to set the property.
- "HOW TO: Remove File Encryption in Windows XP" (308993) tells how to decrypt a file by removing the file encryption property.

Sharing Encrypted Files

The GUI for sharing encrypted files is available only in Windows XP and Windows Server 2003.

- "HOW TO: Share Access to an Encrypted File in Windows XP" (308991) introduces the methodology by which encrypted files can be shared. You can find a short description, including screen shots, in the "5-Minute Security Advisor—Using the Encrypting File System." Remember: sharing encrypted files is a facility that only Windows XP and Windows Server 2003 have.

Top Of Page

Planning for and Recovering Encrypted Files: Recovery Policy

A recovery policy can be an organization's security policy instituted to plan for proper recovery of encrypted files. It's also the policy enforced by Local Security Policy Public Key Policy or Group Policy Public Key Policy. In the latter, the recovery policy specifies how encrypted files may be recovered should the user private key be damaged or lost and the encrypted file unharmed. Recovery certificate(s) are specified in the policy. Recovery can be either data recovery (Windows 2000, Windows XP Professional, and Windows Server 2003) or key recovery (Windows Server 2003 with Certificate Services). Windows 2000 EFS requires the presence of a recovery agent (no recovery agent, no file encryption), but Windows XP and Windows Server 2003 don't. By default, Windows 2000 and Windows Server 2003 have default recovery agents assigned. Windows XP Professional doesn't.

The data recovery process is simple. The user account bound to the recovery agent certificate is used to decrypt the file. The file should then be delivered in a secure manner to the file owner, who may then encrypt the file. Recovery via automatically archived keys is available only with Windows Server 2003 Certificate Services. Additional configuration beyond the installation of Certificate Services is required. In either case, it's most important that a written policy and procedures for recovery are in place. These procedures, if well written and if followed, can ensure that recovery keys and agents are available for use and that recovery is securely carried out. Keep in mind that there are two definitions for "recovery policy." The first definition refers to a written recovery policy and procedures that describe the who, what, where, and when of recovery, as well as what steps should be taken to ensure recovery components are available. The second definition, which is often referred to in the documents below, is the Public Key Policy that's part of the Local Security Policy on stand-alone systems, or Group Policy in a domain. It can specify which certificates are used for recovery, as well as other aspects of Public Key Policies in the domain. You can find more information in the following documents:

- Windows XP and Windows Server 2003 documentation includes steps "To Add a Recovery Agent for a Domain."

- A "Five-Minute Security Advisor—Recovering Encrypted Data Using EFS" explains the importance of backing up encrypted files and EFS keys as well as the basics of recovery.
- "HOW TO: Back Up the Recovery Agent Encrypting File System Private Key in Windows 2000" (241201) explains how archiving the private key of the recovery agent ensures that it will be available to recover EFS files that are protected by it.
- User and recovery agent private keys should be archived.
- If the recovery private key is corrupt or lost, you can create a new Enterprise Data Recovery Policy in Windows 2000. Use the article "HOW TO: Reinitialize the EDRP on a Workgroup Computer Running Windows 2000" (257705) to do so. However, you should realize that this won't allow you to recover previously encrypted files. If a backup of the previous recovery agent certificate and private key is available, those keys should be used. If a new policy is implemented, currently encrypted files should be decrypted and re-encrypted so that the new recovery agent certificate can be used and thus the files will be recoverable.
- Information on the existence of a recovery agent under the control of an administrator is mentioned in" Methods for Recovering Encrypted Data Files"(255742).
- Instructions on using Ntbackup to back up encrypted files, as well as information on system configuration and how to use Ntbackup to restore these files, are discussed in "HOW TO: Use Ntbackup to Recover an Encrypted File or Folder in Windows 2000" (313277).
- The first step in recovery is determining the recovery agent. "Using Efsinfo.exe to Determine Information About Encrypted Files" (243026) describes how to do this using the Windows 2000 Resource Kit tool, esfinfo.exe. The Advanced file properties of encrypted files in Windows XP and Windows Server 2003 display this information automatically.
- "The Local Administrator Is Not Always the Default Encrypting File System Recovery Agent" (255026) explains why the first account defined (during installation) in a Windows 2000 Professional computer becomes the recovery agent.

Disabling or Preventing Encryption

You may decide that you don't wish users to have the ability to encrypt files. By default, they do. You may decide that specific folders shouldn't contain encrypted files. You may also decide to disable EFS until you can implement a sound EFS policy and train users in proper procedures. There are different ways of disabling EFS depending on the operating system and the desired effect:

- System folders cannot be marked for encryption. EFS keys aren't available during the boot process; thus, if system files were encrypted, the system file couldn't boot. To prevent other folders being marked for encryption, you can mark them as system folders. If this isn't possible, then a method to prevent encryption within a folder is defined in "Encrypting File System."
- NT 4.0 doesn't have the ability to use EFS. If you need to disable EFS for Windows 2000 computers joined to a Windows NT 4.0 domain, see "Need to Turn Off EFS on a Windows 2000-Based Computer in Windows NT 4.0-Based Domain" (288579). The registry key mentioned can also be used to disable EFS in Window XP Professional and Windows Server 2003.
- Disabling EFS for Windows XP Professional can also be done by clearing the checkbox for the property page of the Local Security Policy Public Key Policy. EFS can be disabled in XP and Windows Server 2003 computers joined in a Windows Server 2003 domain by clearing the checkbox for the property pages of the domain or organizational unit (OU) Group Policy Public Key Policy.

- "HOW TO: Disable/Enable EFS on a Stand-Alone Windows 2000-Based Computer" (243035) details how to save the recovery agent's certificate and keys when disabling EFS so that you can enable EFS at a future date.
- "HOW TO: Disable EFS for All Computers in a Windows 2000-Based Domain" (222022) provides the best instruction set and clearly defines the difference between deleted domain policy (an OU-based policy or Local Security Policy can exist) versus Initialize Empty Policy (no Windows 2000 EFS encryption is possible throughout the domain).

Special Operations

Let enough people look at anything, and you'll find there are questions that are just not answered by existing documentation or options. A number of these issues, third-party considerations, and post introduction issues can be resolved by reviewing the following articles.

- Specifications for the use of a third-party Certification Authority (CA) can be found at "Third-Party Certification Authority Support for Encrypting File System" (273856). If you wish to use third-party CA certificates for EFS, you should also investigate certificate revocation processing. Windows 2000 EFS certificates aren't checked for revocation. Windows XP and Windows Server 2003 EFS certificates are checked for revocation in some cases, and third-party certificates may be rejected. Information about certificate revocation handling in EFS can be found in the white paper "Encrypting File System in Windows XP and Windows Server 2003".
- When an existing plaintext file is marked for encryption, it's first copied to a temporary file. When the process is complete, the temporary file is marked for deletion, which means portions of the original file may remain on the disk and could potentially be accessible via a disk editor. These bits of data, referred to as *data shreds* or *remanence,* may be permanently removed by using a revised version of the cipher.exe tool. The tool is part of Service Pack 3 (SP3) for Windows 2000 and is included in Windows Server 2003. Instructions for using the tool, along with the location of a downloadable version, can be found in "HOW TO: Use Cipher.exe to Overwrite Deleted Data in Windows" (315672) and in "Cipher.exe Security Tool for the Encrypting File System" (298009).
- How to make encrypted files display in green in Windows Explorer is explained in "HOW TO: Identify Encrypted Files in Windows XP" (320166).
- "How to Enable the Encryption Command on the Shortcut Menu" (241121) provides a registry key to modify for this purpose.
- You may wish to protect printer spool files or hard copies of encrypted files while they're printing. Encryption is transparent to the printing process. If you have the right (possess the key) to decrypt the file and a method exists for printing files, the file will print. However, two issues should concern you. First, if the file is sensitive enough to encrypt, how will you protect the printed copy? Second, the spool file resides in the <system root>\system32\Spool\Printers folder. How can you protect it while its there? You could encrypt that folder, but that would slow printing enormously. The Windows 2000 Resource Kit proposes a separate printer for the printing of these files and how to best secure that printer in the Distributed Systems, Distributed Security, Encrypting Files System, Printing EFS Files section.

Top Of Page

How EFS Works

EFS Architecture and Algorithms

To understand EFS, and therefore anticipate problems, envision potential attacks, and troubleshoot and protect EFS-encrypted files, you should understand the architecture of EFS and the basic encryption, decryption, and recovery algorithms. Much of this information is in the Windows 2000 Resource Kit Distributed Systems Guide, the Windows XP Professional Resource Kit, and the white paper, "Encrypting File System in Windows XP and Windows Server 2003." Many of the algorithms are also described in product documentation. The examples that follow are from the Windows XP Professional Resource Kit:

- A straightforward discussion of the components of EFS, including the EFS service, EFS driver, and the File System Run Time Library, is found in "Components of EFS," a subsection of Chapter 17, "Encrypting File System" in the Windows XP Professional Resource Kit.
- A description of the encryption, decryption, and recovery algorithms EFS uses is in the Resource Kit section "How Files Are Encrypted." This section includes a discussion of the file encryption keys (FEKs) and file Data Recovery Fields and Data Decryption Fields used to hold FEKs encrypted by user and recovery agent public keys.
- "Working with Encryption" includes how-to steps that define the effect of decisions made about changing the encryption properties of folders. The table defines what happens for each file (present, added later, or copied to the folder) for the choice "This folder only" or the option "This folder, subfolders and files."
- "Remote EFS Operations on File Shares and Web Folders" defines what happens to encrypted files and how to enable remote storage.

Top Of Page

Key Differences Between EFS on Windows 2000, Windows XP, and Windows Server 2003

EFS was introduced in Windows 2000. However, there are differences when compared with Windows XP Professional EFS and Windows Server 2003 EFS, including the following:

- You can authorize additional users to access encrypted files (see the section "Sharing Encrypted Files", above). In Windows 2000, you can implement a programmatic solution for the sharing of encrypted files; however, no interface is available. Windows XP and Windows Server 2003 have this interface.
- Offline files can be encrypted. See "HOW TO: Encrypt Offline Files to Secure Data in Windows XP."
- Data recovery agents are recommended but optional. XP doesn't automatically include a default recovery agent. XP will take advantage of an existing Windows 2000 domain-level recovery agent if one is present, but the lack of a domain recovery agent won’t prevent encryption of files on an XP system. A self-signed recovery agent certificate can be requested by using the cipher /R:*filename* command, where *filename* is the name that will be used to create a *.cer file to hold the certificate and a *.pfx file to hold the certificate and private key.
- The Triple DES (3DES) encryption algorithm can be used to replace Data Encryption Standard X (DESX), and after XP SP1, Advanced Encryption Standard (AES) becomes the default encryption algorithm for EFS.
- For Windows XP and Windows Server 2003 local accounts, a password reset disk can be used to safely reset a user's password. (Domain passwords cannot be reset using the disk.) If an administrator uses the "reset password" option from the user's account in the Computer Management console users container, EFS files won't be accessible. If users change the password back to the previous password, they can regain access to encrypted files. To create a password reset disk and for instructions about

how to use a password reset disk, see product documentation and/or the article "HOW TO: Create and Use a Password Reset Disk for a Computer That Is Not a Domain Member in Windows XP" (305478).

- Encrypted files can be stored in Web folders. The Windows XP Professional Resource Kit section "Remote EFS Operations in a Web Folder Environment" explains how.

Windows Server 2003 incorporates the changes introduced in Windows XP Professional and adds the following:

- A default domain Public Key recovery policy is created, and a recovery agent certificate is issued to the Administrator account.
- Certificate Services include the ability for customization of certificate templates and key archival. With appropriate configuration, archival of user EFS keys can be instituted and recovery of EFS-encrypted files can be accomplished by recovering the user's encryption keys instead of decrypting via a file recovery agent. A walk-through providing a step-by-step configuration of Certificate Services for key archival is available in "Certificate Services Example Implementation: Key Archival and Recovery."
- Windows Server 2003 enables users to back up their EFS key(s) directly from the command line and from the details property page by clicking a "Backup Keys" button.

Top Of Page

Misuse and Abuse of EFS and How to Avoid Data Loss or Exposure

Unauthorized persons may attempt to obtain the information encrypted by EFS. Sensitive data may also be inadvertently exposed. Two possible causes of data loss or exposure are misuse (improper use of EFS) or abuse (attacks mounted against EFS-encrypted files or systems where EFS-encrypted files exist).

Inadvertent Problems Due to Misuse

Several issues can cause problems when using EFS. First, when improperly used, sensitive files may be inadvertently exposed. In many cases this is due to improper or weak security policies and a failure to understand EFS. The problem is made all the worse because users *think* their data is secure and thus may not follow usual precautionary methods. This can occur in several scenarios:

- If, for example, users copy encrypted files to FAT volumes, the files will be decrypted and thus no longer protected. Because the user has the right to decrypt files that they encrypted, the file is decrypted and stored in plaintext on the FAT volume. Windows 2000 gives no warning when this happens, but Windows XP and Windows Server 2003 do provide a warning.
- If users provide others with their passwords, these people can log on using these credentials and decrypt the user's encrypted files. (Once a user has successfully logged on, they can decrypt any files the user account has the right to decrypt.)
- If the recovery agent's private key isn't archived and removed from the recovery agent profile, any user who knows the recovery agent credentials can log on and transparently decrypt any encrypted files.

By far, the most frequent problem with EFS occurs when EFS encryption keys and/or recovery keys aren't archived. If keys aren't backed up, they cannot be replaced when lost. If keys cannot be used or replaced, data can be lost. If Windows is reinstalled (perhaps as the result of a disk crash) the keys are destroyed. If a user's profile is damaged, then keys are destroyed. In these, or in any other cases in which keys are damaged or lost and backup keys are unavailable, then encrypted files cannot be decrypted. The encryption keys are bound to the user account, and a new iteration of the operating system means new user accounts. A new user profile means new user keys. If keys are archived, or exported, they can be imported to a new account. If a revocation agent for the files exists, then that account can be used to recover the files. However, in many cases in which keys are destroyed, both user and revocation keys are absent and there is no backup, resulting in lost data.

Additionally, many other smaller things may render encrypted files unusable or expose some sensitive data, such as the following:

- "EFS Files Appear Corrupted When You Open Them" (329741) explains that AES is used to encrypt files after XP SP1 has been installed. This means that these files cannot be decrypted if they're moved to a pre-XP SP1 computer or a Windows 2000 computer since the AES algorithm won't be available.
- "EFS, Credentials, and Private Keys from Certificates Are Unavailable After a Password Is Reset" (290260).
- "User Cannot Gain Access to EFS Encrypted Files After Password Change or When Using a Roaming Profile."
- You can find instructions for using cipher.exe at "HOW TO: Use Cipher.exe to Overwrite Deleted Data in Windows" (315672), which introduces this new tool.
- "Access Is Denied Error Message When Encrypting or Decrypting Files or Folders" (264064) may be the result when encrypting or attempting to encrypt system folders.
- Don't encrypt system files. "Logon Process Hangs After Encrypting Files on Windows 2000," (269397) for example, explains that if you've encrypted a system file such as Autoexec.bat, the file cannot be decrypted because its processed before logon.

Finally, keeping data secure takes more than simply encrypting files. A systems-wide approach to security is necessary. You can find several articles that address best practices for systems security on the TechNet Best Practices page at http://www.microsoft.com/technet/archive/security/bestprac/bpent/sec2/secentbb.mspx. The articles include

- "Security Considerations for End Systems"
- "Security Considerations for Administrative Authority" discusses security in an enterprise
- "Security Entities Building Block Architecture"

Attacks and Countermeasures: Additional Protection Mechanisms for Encrypted Files

Any user of encrypted files should recognize potential weaknesses and avenues of attack. Just as its not enough to lock the front door of a house without considering back doors and windows as avenues for a burglar, encrypting files alone isn't enough to ensure confidentiality.

- Use defense in depth and use file permissions. The use of EFS doesn't obviate the need to use file permissions to limit access to files. File permissions should be used in addition to EFS. If users have obtained encryption keys, they can import them to their account and decrypt files. However, if the user accounts are denied access to the file, the users will be foiled in their attempts to gain this sensitive information.
- Use file permissions to deny delete. Encrypted files can be deleted. If attackers cannot decrypt the file, they may choose to simply delete it. While they don't have the sensitive information, you don't have your file.
- Protect user credentials. If an attacker can discover the identity and password of a user who can decrypt a file, the attacker can log on as that user and view the files. Protecting these credentials is paramount. A strong password policy, user training on devising strong passwords, and best practices on protecting these credentials will assist in preventing this type of attack. An excellent best practices approach to password policy can be found in the Windows Server 2003 product documentation. If account passwords are compromised, anyone can log on using the user ID and password. Once user

have successfully logged on, they can decrypt any files the user account has the right to decrypt. The best defense is a strong password policy, user education, and the use of sound security practices.

- Protect recovery agent credentials. Similarly, if an attacker can log on as a recovery agent, and the recovery agent private key hasn't been removed, the attacker can read the files. Best practices dictate the removal of the recovery agent keys, the restriction of this account's usage to recovery work only, and the careful protection of credentials, among other recovery policies. The sections about recovery and best practices detail these steps.
- Seek out and manage areas where plaintext copies of the encrypted files or parts of the encrypted files may exist. If attackers have possession of, or access to, the computer on which encrypted files reside, they may be able to recover sensitive data from these areas, including the following:
  - Data shreds (remanence) that exist after encrypting a previously unencrypted file (see the "Special Operations" section of this paper for information about using cipher.exe to remove them)
  - The paging file (see "Increasing Security for Open Encrypted Files," an article in the Windows XP Professional Resource Kit, for instructions and additional information about how to clear the paging file on shutdown)
  - Hibernation files (see "Increasing Security for Open Encrypted Files" at http://technet.microsoft.com/library/bb457116.aspx)
  - Temporary files (to determine where applications store temporary files and encrypt these folders as well to resolve this issue
  - Printer spool files (see the "Special Operations" section)
- Provide additional protection by using the System Key. Using Syskey provides additional protection for password values and values protected in the Local Security Authority (LSA) Secrets (such as the master key used to protect user's cryptographic keys). Read the article "Using the System Key" in the Windows 2000 Resource Kit's Encrypting File System chapter. A discussion of the use of Syskey, and possible attacks against a Syskey-protected Windows 2000 computer and countermeasures, can be found in the article "Analysis of Alleged Vulnerability in Windows 2000 Syskey and the Encrypting File System."

Top Of Page

Remote Storage of Encrypted Files Using SMB File Shares and WebDAV

If your policy is to require that data is stored on file servers, not on desktop systems, you will need to choose a strategy for doing so. Two possibilities exist—either storage in normal shared folders on file servers or the use of web folders. Both methods require configuration, and you should understand their benefits and risks.

- If encrypted files are going to be stored on a remote server, the server must be configured to do so, and an alternative method, such as IP Security (IPSec) or Secure Sockets Layer (SSL), should be used to protect the files during transport. Instructions for configuring the server are discussed in "Recovery of Encrypted Files on a Server" (283223) and "HOW TO: Encrypt Files and Folders on a Remote Windows 2000 Server" (320044). However, the latter doesn't mention a critical step, which is that the remote server must be trusted for delegation in Active Directory. Quite a number of articles can be found, in fact, that leave out this step. If the server isn't trusted for delegation in Active Directory, and a user attempts to save the file to the remote server, an "Access Denied" error message will be the result.
- If you need to store encrypted files on a remote server in plaintext (local copies are kept encrypted), you can. The server must, however, be configured to make this happen. You should also realize that

once the server is so configured, no encrypted files can be stored on it. See the article "HOW TO: Prevent Files from Being Encrypted When Copied to a Server" (302093).

- You can store encrypted files in Web folders when using Windows XP or Windows Server 2003. The Windows XP Professional Resource Kit section "Remote EFS Operations in a Web Folder Environment" explains how.
- If your Web applications need to require authentication to access EFS files stored in a Web folder, the code for using a Web folder to store EFS files and require authentication to access them is detailed in "HOW TO: Use Encrypting File System (EFS) with Internet Information Services" (243756).

Top Of Page

Best Practices for SOHO and Small Businesses

Once you know the facts about EFS and have decided how you are going to use it, you should use these documents as a checklist to determine that you have designed the best solution.

- "Best Practices for Encrypting File System" (223316) lists several best practices.
- Best Practices: Windows 2000 Resource Kit, "Administrative Procedures", an article in the EFS chapter of the Windows 2000 Resource Kit, provides insight into the management procedures that should or can be done, including ensuring recovery, disabling EFS, recovery, configuring the agent policy, and viewing recovery agent information. These are best practices from the administrative perspective.
- "Best Practices for Encrypting File System" is included in the Windows 2000 Server product documentation.
- The white paper "Encrypting File System in Windows XP and Windows Server 2003" provides much information about the pluses and minuses of different EFS techniques and many sound practices for managing encryption.

Top Of Page

Enterprise How-tos

How to Implement Data Recovery Strategies with PKI and How to Implement Key Recovery with PKI

By default, EFS certificates are self-signed; that is, they don't need to obtain certificates from a CA. When a user first encrypts a file, EFS looks for the existence of an EFS certificate. If one isn't found, it looks for the existence of a Microsoft Enterprise CA in the domain. If a CA is found, a certificate is requested from the CA; if it isn't, a self-signed certificate is created and used. However, more granular control of EFS, including EFS certificates and EFS recovery, can be established if a CA is present. You can use Windows 2000 or Windows Server 2003 Certificate Services. The following articles explain how.

- "Using a Certificate Authority for the Encrypting File Service" (223338) provides three reasons for using a CA.
- "Using the Cipher.exe Utility to Migrate Self-Signed Certificates to Certification Authority – Issued Certificates" (295680) explains that using cipher /k will archive the self-signed certificate and request a new EFS certificate from the CA.
- User, EFS, and Administrator certificates support EFS use; recovery agent certificates are required for recovery operation.

- Implementation of certificate services for public key infrastructure (PKI) is detailed in the article "Step-by-Step Guide to Encrypting File System (EFS)" and in "Certificate Services Example Implementation: Key Archival and Recovery".

Top Of Page

Troubleshooting

Troubleshooting EFS is easier if you understand how EFS works. There are also well known causes for many of the common problems that arise. Here are a few common problems and their solutions:

- You changed your user ID and password and can no longer decrypt your files. There are two possible approaches to this problem, depending on what you did. First, if the user account was simply renamed and the password reset, the problem may be that you're using XP and this response is expected. When an administrator resets an XP user's account password, the account's association with the EFS certificate and keys is removed. Changing the password to the previous password can reestablish your ability to decrypt your files. For more information, see "User Cannot Gain Access to EFS Encrypted Files After Password Change or When Using a Roaming Profile" (331333), which explains how XP Professional encrypted files cannot be decrypted, even by the original account, if an administrator has changed the password. Second, if you truly have a completely different account (your account was damaged or accidentally deleted), then you must either import your keys (if you've exported them) or ask an administrator to use recovery agent keys (if implemented) to recover the files. Restoring keys is detailed in "HOW TO: Restore an Encrypting File System Private Key for Encrypted Data Recovery in Windows 2000" (242296). How to use a recovery agent to recover files is covered in "Five-Minute Security Advisor—Recovering Encrypted Data Using EFS."
- You've formatted your hard disk and reinstalled the operating system and cannot decrypt your encrypted files. Unless you've exported your EFS keys, or a recovery agent existed and those keys are available, you may not be able to decrypt your files. If your keys, or those of the recovery agent, are available, then it should be possible to either import your keys and decrypt the file or import the recovery agent keys (if necessary) and recover the file. You can determine who the recovery agent of a file is by using esfinfo.exe in Windows 2000 or by looking at the Advanced file properties in XP Professional or Windows Server 2003.
- There is no Advanced button on the file properties page of your Windows XP Home computer, so you cannot mark the file for encryption. No solution is necessary because Windows XP Home doesn't have EFS.

Many other common issues have to do with why users get "Access Denied" messages. (The reason is that they're attempting to access files encrypted by someone else.) By far, however, the largest issue is the recovery of EFS files after a disk crash. See the following articles for troubleshooting other EFS issues:

- The Troubleshooting EFS section of the Windows 2000 Resource Kit includes information explaining that virus checkers can check only files encrypted by the current users.
- "FIX: Incorrect Text 'There Are No EFS Keys' in Sqlstp.log for Error 6006 During Upgrade" (299494) explains an error message during a SQL Server upgrade that has nothing to do with EFS.
- "Users with Roaming Profiles Cannot Use EFS on Domain Controllers" (311513) explains a fix resolved with Windows 2000 SP3.
- "EFS Files Appear Corrupted When You Open Them" (329741) explains the different encryption algorithms used by Windows XP Professional after SP1 and how data might appear corrupted or even

be lost if the files are opened with an earlier version of the OS. It also details how to use an XP registry entry to control which algorithm is used.

- When a profile gets overwritten, the private key of the EFS key pair may no longer be accessible. This may be why the "EFS Recovery Agent Cannot Export Private Keys" (259732).
- Sysprep shouldn't be used on a production machine. Two articles that cover EFS-related problems are "Sysprep.exe May Re-Enable the Encrypting File System" (294844) and "Unable to Access Encrypted Files After Using Sysprep.exe" (288348), which details how Sysprep changes the SIDs on the Administrator and User accounts.
- "HOW TO: Encrypt a File in Windows XP" (307877) includes error messages and warnings that a user may get when attempting to open files encrypted by another user.
- "'Access Is Denied' Error Message When Encrypting or Decrypting Files or Folders" (264064) explains how to remove the system attribute if this is what's blocking file encrypting and what's really desired. (Remember, however, that system files shouldn't be encrypted. The process defined here is for folders that may have inadvertently had the system attribute checked.)
- Encrypted files can be backed up by using Ntbackup, and they remain encrypted. They cannot be restored to a FAT or FAT32 volume. An attempt to do so results in the error message "'Warning: The Restore Destination Device . . . ' Error Message During Restore" (245044).
- "Logon Process Hangs After Encrypting Files on Windows 2000" (269397) explains that if you've encrypted a system file such as Autoexec.bat, the file cannot be decrypted because its processed before logon.
- "Cannot Open Encrypted Files on a Computer with Multiple Windows 2000 Installations" (256168) explains that booting a different operating system means the user is a different user, and thus the encrypted files cannot be decrypted.
- When encrypting many files at once by marking a folder that contains files and/or folder for encryption, or using the Cipher command, you may receive an "Error Message When Attempting to Encrypt Files or Folders" (227465).
- An "Error Message 'Access Denied' When Starting a Recently Installed Program" (272412) may be the result when the temporary folder is encrypted.
- "Recovery of Encrypted Files on Server" (283223) explains how if there is no roaming profile for an account used to remotely encrypt files on a server, a profile is created and a new key pair and certificate are created. This means the keys are different than those used on the local computer.
- "Cannot Gain Access to Microsoft Encrypted File Systems" (243850) explains that a mandatory profile cannot store keys.
- "You Cannot Access Protected Data After You Change Your Password" (322346) explains how changing the password while not connected to the domain can remove access to the original password that was present when the file was encrypted.

Top Of Page

Radical EFS: Using EFS to Encrypt Databases and Using EFS with Other Microsoft Products

"XGEN: Using Windows 2000 Encrypted File System to Encrypt Mdbdata Folder and Contents" (233400) provides information about using EFS to encrypt the Exchange Server 5.5 database.

Top Of Page

Disaster Recovery

You should plan for EFS disaster recovery as part of your Business Continuity Plan. There are three areas of concern:

- In Windows 2000 networks, use a recovery agent and archive user and recovery agent keys.
- Backup of encrypted files should be part of any best practice. Backup of system state is also important, as it provides possible recovery via restoration of the user profile wherein keys lie.
- In Windows 2000 networks that use Certificate Services to provide EFS certificates, and in Windows Server 2003 networks where EFS certificates are used and key archival is exercised, disaster recovery planning should also have plans for the recovery of Certificate Services. The Windows 2000 Resource Kit article "Disaster Recovery Practices" provides some details.

Top Of Page

Overviews and Larger Articles

EFS documentation is primarily located in product Help files, resource kits, and Knowledge Base articles. In addition, some specific columns exist, such as the 5-minute security pieces mentioned earlier and a few white papers.

The following are some recommended white papers, along with other useful references:

"Encrypting File System in Windows XP and Windows Server 2003" is a large white paper which thoroughly documents how EFS works in these operating systems and is must reading for the IT manager or others responsible for designing and deploying formal EFS policy in organizations which use or will use EFS in these operating systems. There are a large number of step-by-step instructions including screen shots that can be helpful to those new to using and administering EFS. However, the paper presumes knowledge of EFS in Windows 2000 and rarely comments on the differences. Key pieces of this document that are essential reading for all are the following:

- "Importing and Exporting Data Recovery Agent Keys."
- "Data Recovery—Best Practices," which discusses the use of a central recovery workstation. The section also covers the use of EFS and CAs and the use of auto enrollment.
- "Encrypting Offline Files."
- "Clearing Page File at Shutdown."
- "Default Encryption Algorithms."
- "Resetting Local Passwords on Windows XP."
- "EFS with WebDAV Folders."
- "EFS and System Restore," which discusses the effect of the use of XP's System Restore on encrypted files.
- "Data Recovery Versus Key Recovery."

"Data Protection and Recovery in Windows XP" was published prior to "Encrypting File Systems in Windows XP and Windows Server 2003" and includes nothing that's not in the latter.

"Encrypting Files System for Windows 2000" was published pre-release of Windows 2000 and is very outdated. Some misinformation is in this guide; therefore, I recommend reading product documentation, Resource Kit material, and Knowledge Base articles instead.

"Step-by-Step Guide to the Encrypting File System" provides a walk-through for implementing EFS in a Windows 2000 domain and includes instructions about restoring EFS-encrypted files and EFS credentials to another computer and setting up Certificate Services for use with EFS.

"Troubleshooting Certificate Status and Revocation" explains how the certificate revocation list (CRL) is checked to determine whether a certificate has been revoked. It also documents that Windows 2000 EFS does no certificate revocation checking, while Windows XP checks the status of certificates added to a file.

The Windows 2000 Resource Kit includes an entire chapter (Chapter 15) about EFS, as well as many references to EFS. The following are a few of the key, or interesting, gems of information in this document:

- The section "Choosing Security Solutions That Use Public Key Technology" documents that EFS uses the base CSP and CryptoAPI and gives a step-by-step explanation about how EFS works.
- "How Encryption Keys Are Protected" explains the five levels of protection for EFS keys.
- "Certificates" lists the requirements for third-party certificates for interoperability with Windows 2000 EFS.
- "Troubleshooting EFS" lists several issues users may encounter when first attempting to use EFS.

The Windows XP Professional Resource Kit includes a chapter about EFS that gives a solid overview of EFS and details XP-specific EFS functions:

- The section "Remote EFS Operations on File Shares and Web Folder" gives the most comprehensive information about using EFS on remote servers.
- "Considerations for Shared Files" raises issues that should be taken into consideration when sharing encrypted files.
- "Data Recovery Implementation Considerations" discusses data recovery in XP and in Windows 2000 domains. Security considerations of different strategies are described.
- "Configuring Recovery Policy in a Standalone Environment" describes how to use the cipher /R:*filename* command to create a self-signed recovery agent certificate for XP systems that aren't joined in a domain.
- "Strengthening Key and File Security" describes how the Data Protection API generates a master key to protect the users private key and to protect the master key.

Where to Download EFS-specific Tools

You can download the following EFS-specific tools by clicking the appropriate link:

Esfinfo.exe: http://www.microsoft.com/downloads/details.aspx?FamilyID=9c70306d-0ef3-4b0c-ab61-81da208f5c47&displaylang=en.

Cipher.exe: http://support.microsoft.com/default.aspx?scid=kb;en-us;298009

How to remove file or folder encryption in Windows XP

This article describes how to decrypt a file that been encrypted by using Encrypting File System (EFS) in Windows XP.

Encryption is converting data into a format that cannot be read by other users. You can use EFS to automatically encrypt your data when it is stored on the hard disk.

Note Only the user who encrypts a file can recover data that has been encrypted, unless the user specifies a recovery agent before they encrypted the files. To make sure that you can decrypt files in the future, you should always export your certificate and private key and keep them in a safe location.

For more information about how to do this, click the following article number to view the article in the Microsoft Knowledge Base:

241201 How to back up the recovery agent Encrypting File System (EFS) private key in Windows Server 2003, in Windows 2000, and in Windows XP

Back to the top

Advanced methods to decrypt a file or a folder

These methods are intended for advanced computer users. If you are not comfortable with advanced methods, you might want to ask someone for help or contact support. For information about how to contact support, visit the following Microsoft Web site:
http://support.microsoft.com/contactus/

To remove encryption from a file or a folder, use the appropriate method later in this section.

Method 1: Remove encryption from a file

Only the following people can decrypt an encrypted file:

- The user who encrypted the file
- Any user who was designated as a recovery agent before the file was encrypted
- Any user who has the public key or private key for the recovery agent or the user who originally encrypted the file
- Any user who has been granted access to the file

Members of the Administrators group cannot decrypt files unless the person who encrypted the files designated them as recovery agents before encrypting the files.

Note You must be the original user who encrypted the file or a designated recovery agent for the file to follow these steps. If you are not authorized to remove encryption, you receive the following error message:

Error Applying Attributes

An error occurred applying attributes to the file:

Path:\Filename

Access is denied

To remove encryption from a file, follow these steps:

1. Use Windows Explorer to locate the encrypted file that you want to decrypt.
2. Right-click the encrypted file, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. Click to clear the **Encrypt contents to secure data** check box, click **OK**, and then click **OK** again.

Method 2: Remove encryption from a folder

Note You must be the original user who encrypted the folder or a designated recovery agent for the folder to follow these steps. If you are not authorized to remove encryption, you receive the following error message:

Error Applying Attributes

An error occurred applying attributes to the file:

Path:\Filename

Access is denied

To remove encryption from a folder, follow these steps:

1. Use Windows Explorer to locate the encrypted folder that you want to decrypt.
2. Right-click the folder, and then click **Properties**.
3. On the **General** tab, click **Advanced**.
4. Click to clear the **Encrypt contents to secure data** check box, click **OK**, and then click **OK** again.
5. When you are prompted to confirm the attribute change:
    - If you want to decrypt only the folder, click **Apply the changes to this folder only**, and then click **OK**.
    - If you want to decrypt the folder and its subfolders and files, click **Apply changes to this folder, subfolders and files**, and then click **OK**.