Part -I

Information Technology - Technical Interview Questions and answer – Networking

What is an IP address?

Every device connected to the public Internet is assigned a unique number known as an Internet Protocol (IP) address. IP addresses consist of four numbers separated by periods (also called a 'dotted-quad') and look something like 127.0.0.1.

In computer networking, an Internet Protocol (IP) address consists of a numerical identification (logical address) that network management assigns to devices participating in a computer network utilizing the Internet Protocol for communication between its nodes.[1] Although computers store IP addresses as binary numbers, they often display them in more human-readable notations, such as 192.168.100.1 (for IPv4), and 2001:db8:0:1234:0:567:1:1 (for IPv6). The role of the IP address has been characterized as follows: "A name indicates what we seek. An address indicates where it is. A route indicates how to get there." [2]

What is subnet Mask ?

A subnet (short for "subnetwork") is an identifiably separate part of an organization's network. Typically, a subnet may represent all the machines at one geographic location, in one building, or on the same local area network (LAN). Having an organization's network divided into subnets allows it to be connected to the Internet with a single shared network address. Without subnets, an organization could get multiple connections to the Internet, one for each of its physically separate subnetworks, but this would require an unnecessary use of the limited number of network numbers the Internet has to assign. It would also require that Internet routing tables on gateways outside the organization would need to know about and have to manage routing that could and should be handled within an organization.

What is ARP? What is ARP Cache Poisoning?

Address Resolution Protocol (ARP) is a protocol for mapping an Internet Protocol address (IP address) to a physical machine address that is recognized in the local network. For example, in IP Version 4, the most common level of IP in use today, an address is 32 bits long. In an Ethernet local area network, however, addresses for attached devices are 48 bits long. (The physical machine address is also known as a Media Access Control or MAC address.) A table, usually called the ARP cache, is used to maintain a correlation between each MAC address and its corresponding IP address. ARP provides the protocol rules for making this correlation and providing address conversion in both directions.

How ARP Works

When an incoming packet destined for a host machine on a particular local area network arrives at a gateway, the gateway asks the ARP program to find a physical host or MAC address that matches the IP address. The ARP program looks in the ARP cache and, if it finds the address, provides it so that the packet can be converted to the right packet length and format and sent to the machine. If no entry is found for the IP address, ARP broadcasts a request packet in a special format to all the machines on the LAN to see if one machine knows that it has that IP address associated with it. A machine that recognizes the IP address as its own returns a reply so indicating. ARP updates the ARP cache for future reference and then sends the packet to the MAC address that replied.

Since protocol details differ for each type of local area network, there are separate ARP Requests for Comments (RFC) for Ethernet, ATM, Fiber Distributed-Data Interface, HIPPI, and other protocols.

There is a Reverse ARP (RARP) for host machines that don't know their IP address. RARP enables them to request their IP address from the gateway's ARP cache.

RARP (Reverse Address Resolution Protocol) is a protocol by which a physical machine in a local area network can request to learn its IP address from a gateway server's Address Resolution Protocol (ARP) table or cache. A network administrator creates a table in a local area network's gateway router that maps the physical machine (or Media Access Control - MAC address) addresses to corresponding Internet Protocol addresses. When a new machine is set up, its RARP client program requests from the RARP server on the router to be sent its IP address. Assuming that an entry has been set up in the router table, the RARP server will return the IP address to the machine which can store it for future use.

What is a default gateway? What happens if I don't have one?

a gateway is a routing device that knows how to pass traffic between different subnets and networks. A computer will know some routes (a route is the address of each node a packet must go through on the Internet to reach a specific destination), but not the routes to every address on the Internet. It won't even know all the routes on the nearest subnets. A gateway will not have this information either, but will at least know the addresses of other gateways it can hand the traffic off to. Your default gateway is on the same subnet as your computer, and is the gateway your computer relies on when it doesn't know how to route traffic.

The default gateway is typically very similar to your IP address, in that many of the numbers may be the same. However, the default gateway is not your IP address. To see what default gateway you are using, follow the steps below for your operating system.

Can a workstation computer be configured to browse the Internet and yet NOT have a default gateway?

What is a subnet?

In computer networks based on the Internet Protocol Suite, a subnetwork, or subnet, is a portion of the network's computers and network devices that have a common, designated IP address routing prefix (cf. Classless Inter-Domain Routing, CIDR).

A routing prefix is the sequence of leading bits of an IP address that precede the portion of the address used as host identifier (or *rest field* in early Internet terminology).

What is APIPA?

What is Automatic Private IP Addressing (APIPA)**?**

A. Windows 98, 98 SE, Me, and 2000 have an Automatic Private IP Addressing (APIPA) feature that will automatically assign an Internet Protocol address to a computer on which it installed. This occurs when the TCP/IP protocol is installed, set to obtain it's IP address automatically from a Dynamic Host Configuration Protocol server, and when there is no DHCP server present or the DHCP server is not available. The Internet Assigned Numbers Authority (IANA) has reserved private IP addresses in the range of **169.254.0.0 - 169.254.255.255** for Automatic Private IP Addressing.

What is an RFC? Name a few if possible (not necessarily the numbers, just the ideas behind them) What is RFC 1918?

Address Allocation for Private Internets February 1996
capabilities of Internet Service Providers. Efforts are in progress within the community to find long term solutions to both of these problems. Meanwhile it is necessary to revisit address allocation procedures, and their impact on the Internet routing system.

What is CIDR?

Short for ***C****lassless* ***I****nter-****D****omain* ***R****outing*, an IP addressing scheme that replaces the older system based on classes A, B, and C. With CIDR, a single IP address can be used to designate many unique IP addresses. A CIDR IP address looks like a normal IP address except that it ends with a slash followed by a number, called the *IP network prefix.* For example:

172.200.0.0/16

The IP network prefix specifies how many addresses are covered by the CIDR address, with lower numbers covering more addresses. An IP network prefix of */12,* for example, can be used to address 1,048,576 former Class C addresses.

CIDR addresses reduce the size of routing tables and make more IP addresses available within organizations.

CIDR is also called *supernetting*.

You have the following Network ID: 192.115.103.64/27. What is the IP range for your network?

192.115.103.65 to 192.115.103.94

You have the following Network ID: 131.112.0.0. You need at least 500 hosts per network. How many networks can you create? What subnet mask will you use?

No of networks = 128 ,

I User Subnet Mask = 255.255.254.0

You need to view at network traffic. What will you use? Name a few tools ?

Monitoring network traffic tool

How do I know the path that a packet takes to the destination?

use "tracert" command-line

What does the ping 192.168.0.1 -l 1000 -n 100 command do?

The ping command will send roundtrip packets to a destination (other PC, router, printer, etc.) and see how long it takes. The 192.168.0.1 is the destination (which, by the way is a typical default IP address of a router.) The -l 1000 is how big the packet should be in bytes. The default is 32, if the -l parameter is not used. And the -n 100 is saying to send, it 100 times. The default is 4, when this parameter is not used.

What is DHCP? What are the benefits and drawbacks of using it?

DHCP is Dynamic Host Configuration Protocol. In a networked environment it is a method to assign an 'address' to a computer when it boots up.
Benefit: A system administrator need not worry about computers being able to access networked resources

Benefits of using DHCP

DHCP provides the following benefits for administering your TCP/IP-based network:

Safe and reliable configuration

DHCP avoids configuration errors caused by the need to manually type in values at each computer. Also, DHCP helps prevent address conflicts caused by a previously assigned IP address being reused to configure a new computer on the network.

Reduces configuration management

Using DHCP servers can greatly decrease time spent configuring and reconfiguring computers on your network. Servers can be configured to supply a full range of additional configuration values when assigning address leases. These values are assigned using DHCP options.

Also, the DHCP lease renewal process helps assure that where client configurations need to be updated often (such as users with mobile or portable computers who change locations frequently), these changes can be made efficiently and automatically by clients communicating directly with DHCP servers.

nteroperability issues

The following section covers issues that affect the use of the DHCP Server service with other services or network configurations.

Using DNS servers with DHCP

Using Routing and Remote Access servers with DHCP

Multihomed DHCP servers

Describe the steps taken by the client and DHCP server in order to obtain an

IP address. ?

DHCP uses a client-server model. The network administrator establishes one or more DHCP servers that maintain TCP/IP configuration information and provide it to clients. The server database includes the following:

Valid configuration parameters for all clients on the network.

Valid IP addresses maintained in a pool for assignment to clients, plus reserved addresses for manual assignment.

Duration of a lease offered by the server. The lease defines the length of time for which the assigned IP address can be used.

With a DHCP server installed and configured on your network, DHCP-enabled clients can obtain their IP address and related configuration parameters dynamically each time they start and join your network. DHCP servers provide this configuration in the form of an address-lease offer to requesting clients.

What is the DHCPNACK and when do I get one? Name 2 scenarios.?

What does DHCPNACK stand for?
DHCP (Dynamic Host Configuration Protocol) Negative Acknowledgment

What ports are used by DHCP and the DHCP clients?

Requests are on UDP port 68, Server replies on UDP 67

Describe the process of installing a DHCP server in an AD infrastructure. .

Open Windows Components Wizard.

Under **Components**, scroll to and click **Networking Services**.

Click **Details**.

Under **Subcomponents of Networking Services**, click **Dynamic Host Configuration Protocol (DHCP)**, and then click **OK**.

Click **Next**. If prompted, type the full path to the Windows Server 2003 distribution files, and then click **Next**.
Required files are copied to your hard disk.

To authorize a DHCP server in Active Directory

Open DHCP.

In the console tree, click **DHCP**.

On the **Action** menu, click **Manage authorized servers**.
The **Manage Authorized Servers** dialog box appears.

Click **Authorize**.

When prompted, type the name or IP address of the DHCP server to be authorized, and then click **OK**.

What is DHCPINFORM?

DHCPInform is a DHCP message used by DHCP clients to obtain DHCP options. While PPP remote access clients do not use DHCP to obtain IP addresses for the remote access connection, Windows 2000 and Windows 98 remote access clients use the DHCPInform message to obtain DNS server IP addresses, WINS server IP addresses, and a DNS domain name. The DHCPInform message is sent after the IPCP negotiation is concluded.

The DHCPInform message received by the remote access server is then forwarded to a DHCP server. The remote access server forwards DHCPInform messages only if it has been configured with the DHCP Relay Agent

Describe the integration between DHCP and DNS?

Traditionally, DNS and DHCP servers have been configured and managed one at a time. Similarly, changing authorization rights for a particular user on a group of devices has meant visiting each one and making configuration changes. DHCP integration with DNS allows the aggregation of these tasks across devices, enabling a company's network services to scale in step with the growth of network users, devices, and policies, while reducing administrative operations and costs.

This integration provides practical operational efficiencies that lower total cost of ownership. Creating a DHCP network automatically creates an associated DNS zone, for example, reducing the number of tasks required of network administrators. And integration of DNS and DHCP in the same database instance provides unmatched consistency between service and management views of IP address-centric network services data

What options in DHCP do you regularly use for an MS network?

What are User Classes and Vendor Classes in DHCP?

Microsoft Vendor Classes

Class Data	Class Name	Description
MSFT 5.0	Microsoft Windows 2000 options	Class that includes all Windows 2000 DHCP clients.
MSFT 98	Microsoft Windows 98 options	Class that includes all Windows 98 and Microsoft Windows Millennium Edition (Me) DHCP clients.
MSFT	Microsoft options	Class that includes all Windows 98, Windows Me, and Windows 2000 DHCP clients.

User Classes

The following list contains pre-defined user classes that are available in Windows 2000 DHCP server.

Unspecified	Default user class	All DHCP clients that have no user class specified.
RRAS.Microsoft	Default Routing and Remote Access class	All Dial-Up Networking (DUN) clients.
Bootp	Default Bootp class	All Bootp clients

How do I configure a client machine to use a specific User Class?

What is the BOOTP protocol used for, where might you find it in Windows network infrastructure?

In computing, **Bootstrap Protocol**, or *BOOTP*, is a UDP network protocol used by a network client to obtain its IP address automatically. This is usually done during the bootstrap process when a computer is starting up. The BOOTP servers assign the IP address to each client from a pool of addresses. We can find, **Bootstrap Protocol** in DHCP Pool configuration in CSCO Switchers and Router.

DNS zones – describe the differences between the 3 types.

DNS stands for Distributed Name System. A DNS server resolves a name to an IP address, as stated in an earlier answer, but it can also point to multiple IP addresses for load balancing, or for backup servers if one or more is offline or not accepting connections.

Individual organizations may have their own DNS servers for their local Intranet.

Some sites have their own DNS server to switch between subdomains within them. For example, a site such as Blogspot can have subdomains come and go quite frequently. Rather than force every DNS server to update their own databases whenever someone creates a new blog, Blogspot could maintain their own DNS server to resolve names within the blogspot.com domain, e.g., to distinguish between myblog.blogspot.com and yourblog.blogspot.com ... their DNS server would be queried once blogspot.com is resolved, and it would be responsible for resolving myblog vs. yourblog.

, such as the Internet. The following are the three main components of DNS:

- **Domain name space and associated resource records (RRs)** A distributed database of name-related information.
- **DNS Name Servers** Servers that hold the domain name space and RRs, and that answer queries from DNS clients.
- **DNS Resolvers** The facility within a DNS client that contacts DNS name servers and issues name queries to obtain resource record information.

DNS Zones

A DNS server that has complete information for part of the DNS name space is said to be the *authority* for that part of the name space. This authoritative information is organized into units called *zones*, which are the main units of replication in DNS. A zone contains one or more RRs for one or more related DNS domains.

The following are the three DNS zone types implemented in Windows 2000:

Standard Primary Holds the master copy of a zone and can replicate it to secondary zones. All changes to a zone are made on the standard primary.

Standard Secondary Contains a read-only copy of zone information that can provide increased performance and resilience. Information in a primary zone is replicated to the secondary by use of the zone transfer mechanism.

Active Directory-integrated A Microsoft proprietary zone type, where the zone information is held in the Windows 2000 Active Directory (AD) and replicated using AD replication.

DNS record types – describe the most important ones.

DNS Resource Records

What Are Resource Records?

An RR is information related to a DNS domain; for example, the host record defining a host IP address. Each RR will contain a common set of information, as follows:

- **Owner** Indicates the DNS domain in which the resource record is found.
- **TTL** The length of time used by other DNS servers to determine how long to cache information for a record before discarding it. For most RRs, this field is optional. The TTL value is measured in seconds, with a TTL value of 0 indicating that the RR contains volatile data that's not to be cached. As an example, SOA records have a default TTL of 1 hour. This prevents these records from being cached by other DNS servers for a longer period, which would delay the propagation of changes.
- **Class** For most RRs, this field is optional. Where it's used, it contains standard mnemonic text indicating the class of an RR. For example, a class setting of IN indicates the record belongs to the Internet (IN) class. At one time there were multiple classes (such as CH for Chaos Net), but today, only the IN class is used.
- **Type** This required field holds a standard mnemonic text indicating the type for an RR. For example, a mnemonic of A indicates that the RR stores host address information.
- **Record-Specific Data** This is a variable-length field containing information describing the resource. This information's format varies according to the type and class of the RR.

Describe the process of working with an external domain name

If it is not possible for you to configure your internal domain as a subdomain of your external domain, use a stand-alone internal domain. This way, your internal and external domain names are unrelated. For example, an organization that uses the domain name contoso.com for their external namespace uses the name corp.internal for their internal namespace.

The advantage to this approach is that it provides you with a unique internal domain name. The disadvantage is that this configuration requires you to manage two separate namespaces. Also, using a stand-alone internal domain that is unrelated to your external domain might create confusion for users because the namespaces do not reflect a relationship between resources within and outside of your network. In addition, you might have to register two DNS names with an Internet name authority if you want to make the internal domain publicly accessible.

Describe the importance of DNS to AD.

When Microsoft began development on Active Directory, full compatibility with the domain name system (DNS) was a critical priority. Active Directory was built from the ground up not just to be fully compatible with DNS but to be so integrated with it that one cannot exist without the other. Microsoft's direction in this case did not just happen by chance, but because of the central role that DNS plays in Internet name resolution and Microsoft's desire to make its product lines embrace the Internet.

While fully conforming to the standards established for DNS, Active Directory can expand upon the standard feature set of DNS and offer some new capabilities such as AD-Integrated DNS, which greatly eases the administration required for DNS environments. In addition, Active Directory can easily adapt to exist in a foreign DNS environment, such as Unix BIND, as long as the BIND version is 8.2.x or higher.

When Microsoft began development on Active Directory, full compatibility with the domain name system (DNS) was a critical priority. Active Directory was built from the ground up not just to be fully compatible with DNS but to be so integrated with it that one cannot exist without the other. Microsoft's direction in this case did not just happen by chance, but because of the central role that DNS plays in Internet name resolution and Microsoft's desire to make its product lines embrace the Internet.

While fully conforming to the standards established for DNS, Active Directory can expand upon the standard feature set of DNS and offer some new capabilities such as AD-Integrated DNS, which greatly eases the administration required for DNS environments. In addition, Active Directory can easily adapt to exist in a foreign DNS environment, such as Unix BIND, as long as the BIND version is 8.2.x or higher.

Describe a few methods of finding an MX record for a remote domain on the Internet.

What does "Disable Recursion" in DNS mean?

In the Windows 2000/2003 DNS console (dnsmgmt.msc), under a server's **Properties** -> **Forwarders** tab is the setting *Do not use recursion for this domain*. On the **Advanced** tab you will find the confusingly similar option *Disable recursion (also disables forwarders)*.

Recursion refers to the action of a DNS server querying additional DNS servers (e.g. local ISP DNS or the root DNS servers) to resolve queries that it cannot resolve from its own database. So what is the difference between these settings?

The DNS server will attempt to resolve the name locally, then will forward requests to any DNS servers specified as forwarders. If *Do not use recursion for this domain* is enabled, the DNS server will pass the query on to forwarders, but will not recursively query any other DNS servers (e.g. external DNS servers) if the forwarders cannot resolve the query.

If *Disable recursion (also disables forwarders)* is set, the server will attempt to resolve a query from its own database *only*. It will not query any additional servers.

If neither of these options is set, the server will attempt to resolve queries normally:
... the local database is queried
... if an entry is not found, the request is passed to any forwarders that are set
... if no forwarders are set, the server will query servers on the **Root Hints** tab to resolve queries beginning at the root domains.

What could cause the Forwarders and Root Hints to be grayed out?

What is a "Single Label domain name" and what sort of issues can it cause?

Single-label names consist of a single word like "contoso".
• Single-label DNS names cannot be registered by using an Internet registrar.
• Client computers and domain controllers that joined to single-label domains require additional configuration to dynamically register DNS records in single-label DNS zones. • Client computers and domain controllers may require additional configuration to resolve DNS queries in single-label DNS zones.
• By default, Windows Server 2003-based domain members, Windows XP-based domain members, and Windows 2000-based domain members do not perform dynamic updates to single-label DNS zones.
• Some server-based applications are incompatible with single-label domain names. Application support may not exist in the initial release of an application, or support may be dropped in a future release. For example, Microsoft Exchange Server 2007 is not supported in environments in which single-label DNS is used.
• Some server-based applications are incompatible with the domain rename feature that is supported in Windows Server 2003 domain controllers and in Windows Server 2008 domain

controllers. These incompatibilities either block or complicate the use of the domain rename feature when you try to rename a single-label DNS name to a fully qualified domain name.

What is the "in-addr.arpa" zone used for?

In a Domain Name System (DNS) environment, it is common for a user or an application to request a Reverse Lookup of a host name, given the IP address. This article explains this process.

The following is quoted from RFC 1035:
"The Internet uses a special domain to support gateway location and Internet address to host mapping. Other classes may employ a similar strategy in other domains. The intent of this domain is to provide a guaranteed method to perform host address to host name mapping, and to facilitate queries to locate all gateways on a particular network on the Internet.

"The domain begins at IN-ADDR.ARPA and has a substructure which follows the Internet addressing structure.

"Domain names in the IN-ADDR.ARPA domain are defined to have up to four labels in addition to the IN-ADDR.ARPA suffix. Each label represents one octet of an Internet address, and is expressed as a character string for a decimal value in the range 0-255 (with leading zeros omitted except in the case of a zero octet which is represented by a single zero).

"Host addresses are represented by domain names that have all four labels specified."
Reverse Lookup files use the structure specified in RFC 1035. For example, if you have a network which is 150.10.0.0, then the Reverse Lookup file for this network would be 10.150.IN-ADDR.ARPA. Any hosts with IP addresses in the 150.10.0.0 network will have a PTR (or 'Pointer') entry in 10.150.IN- ADDR.ARPA referencing the host name for that IP address. A single IN- ADDR.ARPA file may contain entries for hosts in many domains.

Consider the following scenario. There is a Reverse Lookup file 10.150.IN-ADDR.ARPA with the following contents:

Exp : 1.20 IN PTR WS1.ACME.COM.

What are the requirements from DNS to support AD?

When you install Active Directory on a member server, the member server is promoted to a domain controller. Active Directory uses DNS as the location mechanism for domain controllers, enabling computers on the network to obtain IP addresses of domain controllers.

During the installation of Active Directory, the service (SRV) and address (A) resource records are dynamically registered in DNS, which are necessary for the successful functionality of the domain controller locator (Locator) mechanism.

To find domain controllers in a domain or forest, a client queries DNS for the SRV and A DNS resource records of the domain controller, which provide the client with the names and IP addresses of the domain controllers. In this context, the SRV and A resource records are referred to as Locator DNS resource records.

When adding a domain controller to a forest, you are updating a DNS zone hosted on a DNS server with the Locator DNS resource records and identifying the domain controller. For this reason, the DNS zone must allow dynamic updates (RFC 2136) and the DNS server hosting

that zone must support the SRV resource records (RFC 2782) to advertise the Active Directory directory service. For more information about RFCs, see DNS RFCs.

If the DNS server hosting the authoritative DNS zone is not a server running Windows 2000 or Windows Server 2003, contact your DNS administrator to determine if the DNS server supports the required standards. If the server does not support the required standards, or the authoritative DNS zone cannot be configured to allow dynamic updates, then modification is required to your existing DNS infrastructure.

For more information, see Checklist: Verifying DNS before installing Active Directory and Using the Active Directory Installation Wizard.

Important

The DNS server used to support Active Directory must support SRV resource records for the Locator mechanism to function. For more information, see Managing resource records.

It is recommended that the DNS infrastructure allows dynamic updates of Locator DNS resource records (SRV and A) before installing Active Directory, but your DNS administrator may add these resource records manually after installation. After installing Active Directory, these records can be found on the domain controller in the following location: systemroot\System32\Config\Netlogon.dns

How do you manually create SRV records in DNS?

this is on windows server

go to run ---> dnsmgmt.msc

rightclick on the zone you want to add srv record to and choose "other new record"

and choose service location(srv).....

Name 3 benefits of using AD-integrated zones.

Active Directory–integrated DNS enables Active Directory storage and replication of DNS zone databases. Windows 2000 DNS server, the DNS server that is included with Windows 2000 Server, accommodates storing zone data in Active Directory. When you configure a computer as a DNS server, zones are usually stored as text files on name servers — that is, all of the zones required by DNS are stored in a text file on the server computer. These text files must be synchronized among DNS name servers by using a system that requires a separate replication topology and schedule called a zone transfer However, if you use Active Directory–integrated DNS when you configure a domain controller as a DNS name server, zone data is stored as an Active Directory object and is replicated as part of domain replication.

What are the benefits of using Windows 2003 DNS when using AD-integrated zones?

If your DNS topology includes Active Directory, use Active Directory–integrated zones. Active Directory–integrated zones enable you to store zone data in the Active Directory database. Zone information about any primary DNS server within an Active Directory–integrated zone is always replicated.

Because DNS replication is single-master, a primary DNS server in a standard primary DNS zone can be a single point of failure. In an Active Directory–integrated zone, a primary DNS server cannot be a single point of failure because Active Directory uses multimaster replication. Updates that are made to any domain controller are replicated to all domain controllers and the zone information about any primary DNS server within an Active Directory–integrated zone is always replicated. Active Directory–integrated zones:

- Enable you to secure zones by using secure dynamic update.
- Provide increased fault tolerance. Every Active Directory–integrated zone can be replicated to all domain controllers within the Active Directory domain or forest. All DNS servers running on these domain controllers can act as primary servers for the zone and accept dynamic updates.
- Enable replication that propagates changed data only, compresses replicated data, and reduces network traffic.

If you have an Active Directory infrastructure, you can only use Active Directory–integrated zones on Active Directory domain controllers. If you are using Active Directory–integrated zones, you must decide whether or not to store Active Directory–integrated zones in the application directory partition.

You can combine Active Directory–integrated zones and file-based zones in the same design. For example, if the DNS server that is authoritative for the private root zone is running on an operating system other than Windows Server 2003 or Windows 2000, it cannot act as an Active Directory domain controller. Therefore, you must use file-based zones on that server. However, you can delegate this zone to any domain controller running either Windows Server 2003 or Windows 2000.

You installed a new AD domain and the new (and first) DC has not registered its SRV records in DNS. Name a few possible causes.

What are the benefits and scenarios of using Stub zones?

Understanding stub zones

A stub zone is a copy of a zone that contains only those resource records necessary to identify the authoritative Domain Name System (DNS) servers for that zone. A stub zone is used to resolve names between separate DNS namespaces. This type of resolution may be necessary when a corporate merger requires that the DNS servers for two separate DNS namespaces resolve names for clients in both namespaces.

A stub zone consists of:

- The start of authority (SOA) resource record, name server (NS) resource records, and the glue A resource records for the delegated zone.
- The IP address of one or more master servers that can be used to update the stub zone.

The master servers for a stub zone are one or more DNS servers authoritative for the child zone, usually the DNS server hosting the primary zone for the delegated domain name.

Using stub zones
Updated: January 21, 2005

Using stub zones

Use stub zones to:

- **Keep delegated zone information current.** By updating a stub zone for one of its child zones regularly, the DNS server hosting both the parent zone and the stub zone will maintain a current list of authoritative DNS servers for the child zone.
- **Improve name resolution.** Stub zones enable a DNS server to perform recursion using the stub zone's list of name servers without needing to query the Internet or internal root server for the DNS namespace.
- **Simplify DNS administration.** By using stub zones throughout your DNS infrastructure, you can distribute a list of the authoritative DNS servers for a zone without using secondary zones. However, stub zones do not serve the same purpose as secondary zones and are not an alternative when considering redundancy and load sharing.

There are two lists of DNS servers involved in the loading and maintenance of a stub zone:

- The list of master servers from which the DNS server loads and updates a stub zone. A master server may be a primary or secondary DNS server for the zone. In both cases, it will have a complete list of the DNS servers for the zone.

- The list of the authoritative DNS servers for a zone. This list is contained in the stub zone using name server (NS) resource records.

When a DNS server loads a stub zone, such as widgets.example.com, it queries the master servers, which can be in different locations, for the necessary resource records of the authoritative servers for the zone widgets.example.com. The list of master servers may contain a single server or multiple servers and can be changed anytime. For more information, see Configure a stub zone for local master servers.

What are the benefits and scenarios of using Conditional Forwarding?

Rather than having a DNS server forward all queries it cannot resolve to forwarders, the DNS server can forward queries for different domain names to different DNS servers according to the specific domain names that are contained in the queries. Forwarding according to these domain-name conditions improves conventional forwarding by adding a second condition to the forwarding process.

A conditional forwarder setting consists of a domain name and the IP address of one or more DNS servers. To configure a DNS server for conditional forwarding, a list of domain names is set up on the Windows Server 2003-based DNS server along with the DNS server IP address. When a DNS client or server performs a query operation against a Windows Server 2003-based DNS server that is configured for forwarding, the DNS server looks to see if the query can be resolved by using its own zone data or the zone data that is stored in its cache, and then, if the DNS server is configured to forward for the domain name that is designated in the query (a match), the query is forwarded to the IP address of a DNS Server that is associated with the domain name. If the DNS server has no domain name listed for the name that is designated in the query, it attempts to resolve the query by using standard recursion.

What are the differences between Windows Clustering, Network Load Balancing and Round Robin, and scenarios for each use?

Cluster technologies are becoming increasingly important to ensure service offerings meet the requirements of the enterprise. Windows 2000 and Windows Server 2003 support three cluster technologies to provide high availability, reliability and scalability. These technologies are: NLB, CLB and Server cluster. These technologies have a specific purpose and are designed to meet different requirements.

- **Server cluster** provides failover support for applications and services that require high availability, scalability and reliability, and is ideally suited for back-end applications and services, such as database servers. Server cluster can use various combinations of active and passive nodes to provide failover support for mission critical applications and services.
- **NLB** provides failover support for IP-based applications and services that require high scalability and availability, and is ideally suited for Web tier and front-end

services. NLB clusters can use multiple adapters and different broadcast methods to assist in the load balancing of TCP, UDP and GRE traffic requests.

- **Component Load Balancing** provides dynamic load balancing of middle-tier application components that use COM+ and is ideally suited for application servers. CLB clusters use two clusters. The routing cluster can be configured as a routing list on the front-end Web servers or as separate servers that run Server cluster.

Cluster technologies by themselves are not enough to ensure that high availability goals can be met. Multiple physical locations may be necessary to guard against natural disasters and other events that may cause complete service outage. Effective processes and procedures, in addition to good architecture, are the keys to high availability.

Round robin is a local balancing mechanism used by DNS servers to share and distribute network resource loads. You can use it to rotate all resource record (RR) types contained in a query answer if multiple RRs are found.

By default, DNS uses round robin to rotate the order of RR data returned in query answers where multiple RRs of the same type exist for a queried DNS domain name. This feature provides a simple method for load balancing client use of Web servers and other frequently queried multihomed computers.

If round robin is disabled for a DNS server, the order of the response for these queries is based on a static ordering of RRs in the answer list as they are stored in the zone (either its zone file or Active Directory).

How do I work with the Host name cache on a client computer?

How do I clear the DNS cache on the DNS server?

To clear DNS Cache do the following:

1. Start

2. Run

3. Type "cmd" and press enter

4. In the command window type "ipconfig /flushdns"

5. If done correctly it should say "Successfully flushed the DNS Resolver Cache."

What is the 224.0.1.24 address used for?

WINS server group address. Used to support auto discovery and dynamic configuration of replication for WINS servers. For more information, see WINS replication overview

WINS server group address. Used to support auto discovery and dynamic configuration of replication for WINS servers. For more information, see WINS replication overview

What is WINS and when do we use it?

Microsoft Windows Internet Name Service (WINS) is an RFC-compliant NetBIOS name- to-IP-address mapping service. WINS allows Windows-based clients to easily locate resources on Transmission Control Protocol/Internet Protocol (TCP/IP) networks. WINS servers maintain databases of static and dynamic resource name—to-IP-address mappings. Because the Microsoft WINS database supports dynamic name and IP address entries, WINS can be used with Dynamic Host Configuration Protocol (DHCP) services to provide easy configuration and administration of Windows-based TCP/IP networks.
WINS servers provide the following benefits:

- Dynamic database that supports NetBIOS computer name registration and name resolution in an environment where the dynamic TCP/IP configuration of DHCP-enabled clients is dynamically configured for TCP/IP.
- Centralized management of the NetBIOS computer name database and its replication to other WINS servers.
- Reduction of NetBIOS name query IP broadcast traffic.
- Support for Windows-based clients (including Windows NT Server, Windows NT Workstation, Windows 95, Windows for Workgroups, and LAN Manager 2.*x*).
- Support for transparent browsing across routers for Windows NT Server, Windows NT Workstation, Windows 95, and Windows for Workgroups clients.

to the WINS server. The WINS server returns the destination computer's IP address to the original computer without the need for broadcast traffic.

The second reason for using WINS is that it's dynamic. As computers attach to and detach from the network, the WINS databases are updated automatically. This means that you don't have to create a static LMHOST file that the computers can read to determine IP addresses.

Can you have a Microsoft-based network without any WINS server on it? What are the "considerations" regarding not using WINS?

Describe the differences between WINS push and pull replications.

Microsoft WINS Server Push and Pull Partners

A given network should have one or more WINS servers that WINS clients can contact to resolve a computer name to an IP address. It is desirable to have multiple WINS servers installed on an intranet for the following reasons:

- To distribute the NetBIOS computer name query and registration processing load
- To provide WINS database redundancy, backup, and disaster recovery

Microsoft WINS servers communicate with other Microsoft WINS servers to fully replicate their databases with each other. This ensures that a name registered with one WINS server is replicated to all other Microsoft WINS servers within the intranet, providing a replicated and enterprise-wide database.

When multiple WINS servers are used, each WINS server is configured as a *pull* or *push* partner of at least one other WINS server. The following table describes the pull and push partner types of replication partners.

What is the difference between tombstoning a WINS record and simply deleting it?

Through replication and convergence, the [1C] record ownership will change from WINS server to WINS server. Eventually, you may many end up with a scenario where a WINS server that owns a [1C] record and its direct replication partner has a replica of the [1C] record but does not own the record. The problem occurs when no domain controllers refresh the [1C] record on the remote WINS server, the records will expire, become tombstoned, and be scavenged out of the database. The following is an example of what could happen

Name the NetBIOS names you might expect from a Windows 2003 DC that is registered in WINS.

If a Microsoft Windows NT 3.5-based client computer does not receive a response from the primary Windows Internet Name Service (WINS) server, it queries the secondary WINS server to resolve a NetBIOS name. However, if a NetBIOS name is not found in the primary WINS server's database, a Windows NT 3.5-based client does not query the secondary WINS server.

In Microsoft Windows NT 3.51 and later versions of the Windows operating system, a Windows-based client does query the secondary WINS server if a NetBIOS name is not found in the primary WINS server's database. Clients that are running the following versions
In Windows NT 3.51, Windows NT 4, Windows 95, Windows 98, Windows 2000, Windows Millennium Edition, Windows XP, and Windows Server 2003, you can specify up to 12 WINS servers. Additional WINS servers are useful when a requested name is not found in the primary WINS server's database or in the secondary WINS server's database. In this situation, the WINS client sends a request to the next server in the list.

You can find a list of additional server names in the following registry subkey, where *adapter_guid* represents the GUID of your adapter:

HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services\NetBT\Parameters\Interfaces\Tcpip_<adapter_guid>
Note Make sure that the NameServerList registry entry in this subkey has a multistring type (REG_MULTI_SZ).

What is TCP/IP and Explain some TCP /IP Protocol ?

What is TCP/IP? TCP/IP (Transmission Control Protocol/Internet Protocol) is the basic communication language or protocol of the Internet. It can also be used as a communications protocol in a private network (either an intranet or an extranet). When you are set up with direct access to the Internet, your computer is provided with a copy of the TCP/IP program just as every other computer that you may send messages to or get information from also has a copy of TCP/IP.

TCP/IP is a two-layer program. The higher layer, Transmission Control Protocol, manages the assembling of a message or file into smaller packets that are transmitted over the Internet and received by a TCP layer that reassembles the packets into the original message. The lower layer, Internet Protocol, handles the address part of each packet so that it gets to the right destination. Each gateway computer on the network checks this address to see where to forward the message. Even though some packets from the same message are routed differently than others, they'll be reassembled at the destination.

TCP/IP uses the client/server model of communication in which a computer user (a client) requests and is provided a service (such as sending a Web page) by another computer (a server) in the network. TCP/IP communication is primarily point-to-point, meaning each communication is from one point (or host computer) in the network to another point or host computer. TCP/IP and the higher-level applications that use it are collectively said to be "stateless" because each client request is considered a new request unrelated to any previous one (unlike ordinary phone conversations that require a dedicated connection for the call duration). Being stateless frees network paths so that everyone can use them continuously. (Note that the TCP layer itself is not stateless as far as any one message is concerned. Its connection remains in place until all packets in a message have been received.)

Many Internet users are familiar with the even higher layer application protocols that use TCP/IP to get to the Internet. These include the World Wide Web's Hypertext Transfer Protocol (HTTP), the File Transfer Protocol (FTP), Telnet (Telnet) which lets you logon to remote computers, and the Simple Mail Transfer Protocol (SMTP). These and other protocols are often packaged together with TCP/IP as a "suite."

Personal computer users with an analog phone modem connection to the Internet usually get to the Internet through the Serial Line Internet Protocol (SLIP) or the Point-to-Point Protocol (PPP). These protocols encapsulate the IP packets so that they can be sent over the dial-up phone connection to an access provider's modem.

Protocols related to TCP/IP include the User Datagram Protocol (UDP), which is used instead of TCP for special purposes. Other protocols are used by network host computers for exchanging router information. These include the Internet Control Message Protocol (ICMP), the Interior Gateway Protocol (IGP), the Exterior Gateway Protocol (EGP), and the Border Gateway Protocol (BGP).

What is NetBios ?

Netbios.exe is a NetBIOS programming sample that implements an echo server and client. The sample illustrates how a client and server should be written in order to make the application protocol and LAN Adapter (LANA) independent. It also shows how to avoid common mistakes programmers frequently make when writing NetBIOS applications under WIN32.

Describe the role of the routing table on a host and on a router.

In internetworking, the process of moving a packet of data from source to destination. Routing is usually performed by a dedicated device called a router. Routing is a key feature of the Internet because it enables messages to pass from one computer to another and eventually reach the target machine. Each intermediary computer performs routing by passing along the message to the next computer. Part of this process involves analyzing a *routing table* to determine the best path.

(row´ter) **(n.)** A device that forwards data packets along networks. A router is connected to at least two networks, commonly two LANs or WANs or a LAN and its ISP's network. Routers are located at gateways, the places where two or more networks connect.

Routers use headers and forwarding tables to determine the best path for forwarding the packets, and they use protocols such as ICMP to communicate with each other and configure the best route between any two hosts.

Very little filtering of data is done through routers.

Defined OSI model ?

The '***Open Systems Interconnection Basic Reference Model***' (OSI Reference Model or **OSI Model**) is an abstract description for layered communications and computer network protocol design. It was developed as part of the Open Systems Interconnection (OSI) initiative[1]. In its most basic form, it divides network architecture into seven layers which, from top to bottom, are the Application, Presentation, Session, Transport, Network, Data-Link, and Physical Layers. It is therefore often referred to as the ***OSI Seven Layer Model***.

1 The Physical Layer defines the electrical and physical specifications for devices. In particular, it defines the relationship between a device and a physical medium. This includes the layout of pins, voltages, cable specifications, Hubs, repeaters, network adapters, Host Bus Adapters (HBAs used in Storage Area Networks) and more.

To understand the function of the Physical Layer in contrast to the functions of the Data Link Layer, think of the Physical Layer as concerned primarily with the interaction of a single device with a medium, where the Data Link Layer is concerned more with the interactions of multiple devices (i.e., at least two) with a shared medium. The Physical Layer will tell one device how to transmit to the medium, and another device how to receive from it (in most

cases it does not tell the device how to connect to the medium). Obsolescent Physical Layer standards such as RS-232 do use physical wires to control access to the medium.

2 Data Link Layer

The Data Link Layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the Physical Layer. Originally, this layer was intended for point-to-point and point-to-multipoint media, characteristic of wide area media in the telephone system. Local area network architecture, which included broadcast-capable multiaccess media, was developed independently of the ISO work, in IEEE Project 802. IEEE work assumed sublayering and management functions not required for WAN use. In modern practice, only error detection, not flow control using sliding window, is present in modern data link protocols such as Point-to-Point Protocol (PPP), and, on local area networks, the IEEE 802.2 LLC layer is not used for most protocols on Ethernet, and, on other local area networks, its flow control and acknowledgment mechanisms are rarely used. Sliding window flow control and acknowledgment is used at the Transport Layer by protocols such as TCP, but is still used in niches where X.25 offers performance advantages.

Both WAN and LAN services arrange bits, from the Physical Layer, into logical sequences called frames. Not all Physical Layer bits necessarily go into frames, as some of these bits are purely intended for Physical Layer functions. For example, every fifth bit of the FDDI bit stream is not used by the Data Link Layer.

3 The Network Layer

provides the functional and procedural means of transferring variable length data sequences from a source to a destination via one or more networks, while maintaining the quality of service requested by the Transport Layer. The Network Layer performs network routing functions, and might also perform fragmentation and reassembly, and report delivery errors. Routers operate at this layer—sending data throughout the extended network and making the Internet possible. This is a logical addressing scheme – values are chosen by the network engineer. The addressing scheme is hierarchical.The best-known example of a Layer 3 protocol is the Internet Protocol (IP). It manages the connectionless transfer of data one hop at a time, from end system to ingress router, router to router, and from egress router to destination end system. It is not responsible for reliable delivery to a next hop, but only for the detection of errored packets so they may be discarded. When the medium of the next hop cannot accept a packet in its current length, IP is responsible for **fragmenting** into sufficiently small packets that the medium can accept it.A number of layer management protocols, a function defined in the Management Annex, ISO 7498/4, belong to the Network Layer. These include routing protocols, multicast group management, Network Layer information and error, and Network Layer address assignment. It is the function of the payload that makes these belong to the Network Layer, not the protocol that carries them.

4 Transport Layer

The Transport Layer provides transparent transfer of data between end users, providing reliable data transfer services to the upper layers. The Transport Layer controls the reliability of a given link through flow control, segmentation/desegmentation, and error control. Some protocols are state and connection oriented. This means that the Transport Layer can keep track of the segments and retransmit those that fail.

Although not developed under the OSI Reference Model and not strictly conforming to the OSI definition of the Transport Layer, the best known examples of a Layer 4 protocol are the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

5 The Session Layer

controls the dialogues/connections (sessions) between computers. It establishes, manages and terminates the connections between the local and remote application. It provides for full-duplex, half-duplex, or simplex operation, and establishes checkpointing, adjournment, termination, and restart procedures. The OSI model made this layer responsible for "graceful close" of sessions, which is a property of TCP, and also for session checkpointing and recovery, which is not usually used in the Internet Protocol Suite. The Session Layer is commonly implemented explicitly in application environments that use remote procedure calls (RPCs).

Layer 6: Presentation Layer

The Presentation Layer establishes a context between Application Layer entities, in which the higher-layer entities can use different syntax and semantics, as long as the Presentation Service understands both and the mapping between them. The presentation service data units are then encapsulated into Session Protocol Data Units, and moved down the stack.

The original presentation structure used the Basic Encoding Rules of Abstract Syntax Notation One (ASN.1), with capabilities such as converting an EBCDIC-coded text file to an ASCII-coded file, or serializing objects and other data structures into and out of XML. ASN.1 has a set of cryptographic encoding rules that allows end-to-end encryption between application entities.

7 Application Layer

The application layer is the OSI layer closest to the end user, which means that both the OSI application layer and the user interact directly with the software application. This layer interacts with software applications that implement a communicating component. Such application programs fall outside the scope of the OSI model. Application layer functions typically include identifying communication partners, determining resource availability, and synchronizing communication. When identifying communication partners, the application layer determines the identity and availability of communication partners for an application with data to transmit. When determining resource availability, the application layer must decide whether sufficient network resources for the requested communication exist. In synchronizing communication, all communication between applications requires cooperation that is managed by the application layer. Some examples of application layer implementations include Telnet, File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP).

What are routing protocols? Why do we need them? Name a few.

routing protocol is a protocol that specifies how routers communicate with each other to disseminate information that allows them to select routes between any two nodes on a network. Typically, each router has *a prior* knowledge only of its immediate neighbors. A routing protocol shares this information so that routers have knowledge of the network topology at large. For a discussion of the concepts behind routing protocols, see: Routing.

The term **routing protocol** may refer more specifically to a protocol operating at Layer 3 of the OSI model which similarly disseminates topology information between routers.

Many routing protocols used in the public Internet are defined in documents called RFCs.[1][2][3][4]

There are three major types of routing protocols, some with variants: link-state routing protocols, path vector protocols and distance vector routing protocols.

The specific characteristics of routing protocols include the manner in which they either prevent routing loops from forming or break routing loops if they do form, and the manner in which they determine preferred routes from a sequence of hop costs and other preference factors.

- IGRP (Interior Gateway Routing Protocol)
- EIGRP (Enhanced Interior Gateway Routing Protocol)
- OSPF (Open Shortest Path First)
- RIP (Routing Information Protocol)
- IS-IS (Intermediate System to Intermediate System)

What are router interfaces? What types can they be?

The interfaces on a router provide network connectivity to the router. The console and auxiliary ports are used for managing the router. Routers also have ports for LAN and WAN connectivity.

The LAN interfaces usually include Ethernet, Fast Ethernet, Fiber Distributed Data Interface (FDDI), or Token Ring. The AUI port is used to provide LAN connectivity. You can use a converter to attach your LAN to the router. Some higher-end routers have separate interfaces for ATM (Asynchronous Transfer Mode) as well.

Sync and Async serial interfaces are used for WAN connectivity. ISDN (Integrated Services Digital Network) interfaces are used to provide the ISDN connectivity. Using ISDN, you can transmit both voice and data.

Bas Topology

To prevent collisions senses multi access /collision detection CSMA/CD is used in Ethernet .one way transferring data .

Ethernet is one of the earliest LAN technologies. An Ethernet LAN typically uses special grades of twisted pair cabling. Ethernet networks can also use coaxial cable, but this cable medium is becoming less common. The most commonly installed Ethernet systems are called 10BaseT. The router provides the interfaces for twisted pair cables. A converter can be attached to the AUI port of a router to connect to a 10base2, 10baseT, or 10base5 LAN interface. Ethernet and Token Ring use MAC addressing (physical addressing).

The Ethernet interfaces on the router are E0, E1, E2, and so on. E stands for Ethernet, and the number that follows represents the port number. These interfaces provide connectivity to an Ethernet LAN. In a non-modular Cisco router, the Ethernet ports are named as above, but in modular routers they are named as E0/1, where E stands for Ethernet, 0 stands for slot number, and 1 stands for port number in that slot.

Token Ring Topology

Token Ring is the second most widely used LAN technology after Ethernet, where all computers are connected in a logical ring topology. Physically, each host attaches to an MSAU (Multistation Access Unit) in a star configuration. MSAU's can be chained together to maintain the logical ring topology. An empty frame called a token is passed around the network. A device on the network can transmit data only when the empty token reaches the device. This eliminates collisions on a Token Ring network. Token Ring uses MAC addresses just like any other LAN technology.

The Token Ring interfaces on a non-modular router are To0, To1, To2 and so on. "To" stands for Token Ring and the number following "To" signifies the port number. In a modular router, "To" will be followed by the slot number/port number

FDDI

Fiber Distributed Data Interface (FDDI) is a LAN technology that uses fiber optic cable. FDDI is a ring topology that uses four-bit symbols rather than eight-bit octets in its frames. The 48-bit MAC addresses have 12 four-bit symbols for FDDI. FDDI is very fast and provides a data transfer rate of 100 Mbps and uses a token-passing mechanism to prevent collisions.
FDDI uses two rings with their tokens moving in opposite directions to provide redundancy to the network. Usually only one ring is active at a given time. If one ring breaks, the other ring is used and the network does not experience downtime.
FDDI interfaces on a non-modular Cisco router are F0, F1, F2 and so on. "F" stands for FDDI and the number following "F" signifies the port number. In a modular router, a slot number/port number will follow "F".

ISDN

Integrated Services Digital Network (ISDN) is a set of ITU-T (Telecommunication Standardization Sector of the International Telecommunications Union) standards for digital transmission over ordinary telephone copper wire as well as over other media. ISDN provides the integration of both analog or voice data together with digital data over the same network. ISDN has two levels of service:

- Basic Rate Interface (BRI)
- Primary Rate Interface (PRI)

The BRI interfaces for ISDN on a non-modular router are BRI0, BRI1, and so on, with the number following "BRI" signifying the port number. In a modular router, BRI is followed by the slot number/port number.

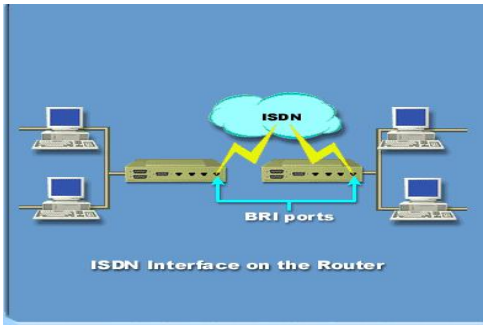

Synchronous transmission signals

occur at the same clock rate and all clocks are based on a single reference clock. Since asynchronous transmission is a character-by-character transmission type, each character is delimited by a start and stop bit, therefore clocks are not needed in this type of transmission. Synchronous communication requires a response at the end of each exchange of frames, while asynchronous communications do not require responses.

Support for the Synchronous Serial interface is supplied on the Multiport Communications Interface (CSC-MCI) and the Serial Port Communications Interface (CSC-SCI) network interface cards. The Asynchronous Serial interface is provided by a number of methods, including RJ-11, RJ-45, and 50-pin Telco connectors.

Some ports can function both as Synchronous Serial interfaces and Asynchronous Serial interfaces. Such ports are called Async/Sync ports. The Async/Sync ports support Telco and RJ-11 connectors.

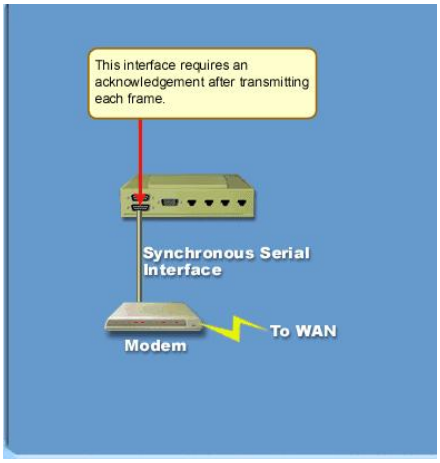

In Windows 2003 routing, what are the interface filters?

What is NAT?

This article will describe how to setup and configure NAT in Windows 2003. NAT, or Network Address Translation, is a widely used IP translation and mapping protocol that works on the network layer (level 3) of the OSI model. It is sometimes referred to as a routing

protocol because of the way it allows packets from a private network to be routed to the Internet

NAT actsas a middle man between the internal and external network; packets coming from the private network are handled by NAT and then transferred to their intended destination.

A single external address is used on the Internet so that the internal IP addresses are not shown. A table is created on the router that lists local and global addresses and uses it as a reference when translating IP addresses.

This article will describe how to setup and configure NAT in Windows 2003. NAT, or Network Address Translation, is a widely used IP translation and mapping protocol that works on the network layer (level 3) of the OSI model. It is sometimes referred to as a routing protocol because of the way it allows packets from a private network to be routed to the Internet

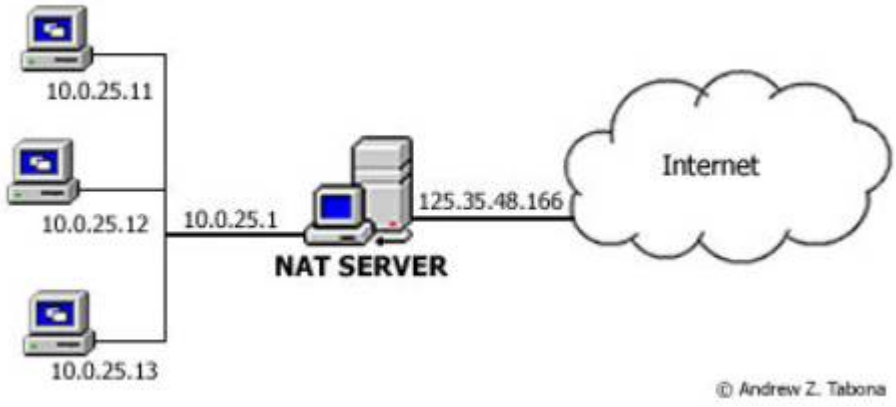

NAT can work in several ways:

Static NAT

An unregistered IP address is mapped to a registered IP address on a one-to-one basis - which is useful when a device needs to be accessed from outside the network.

Dynamic NAT

An unregistered IP address is mapped to a registered IP address from a group of registered IP addresses. For example, a computer 192.168.10.121 will translate to the first available IP in a range from 212.156.98.100 to 212.156.98.150.

Overloading

A form of dynamic NAT, it maps multiple unregistered IP addresses to a single registered IP address, but in this case uses different ports. For example, IP address 192.168.10.121 will be mapped to 212.56.128.122:port_number (212.56.128.122:1080).

Overlapping

This when addresses in the inside network overlap with addresses in the outside network - the IP addresses are registered on another network too. The router must maintain a lookup table of these addresses so that it can intercept them and replace them with registered unique IP addresses.

How NAT works

A table of information about each packet that passes through is maintained by NAT.

When a computer on the network attempts to connect to a website on the Internet:

- the header of the source IP address is changed and replaced with the IP address of the NAT computer on the way out
- the "destination" IP address is changed (based on the records in the table) back to the specific internal private class IP address in order to reach the computer on the local network on the way back in

Network Address Translation can be used as a basic firewall – the administrator is able to filter out packets to/from certain IP addresses and allow/disallow access to specified ports. It is also a means of saving IP addresses by having one IP address represent a group of computers.

Setting up NAT

To setup NAT you must start by opening the Configure your server wizard in administrative tools and selecting the RRAS/VPN Server role. Now press next and the RRAS setup wizard will open. The screen below shows the Internet Connection screen in which you must specify which type of connection to the Internet and whether or not you want the basic firewall feature to be enabled.

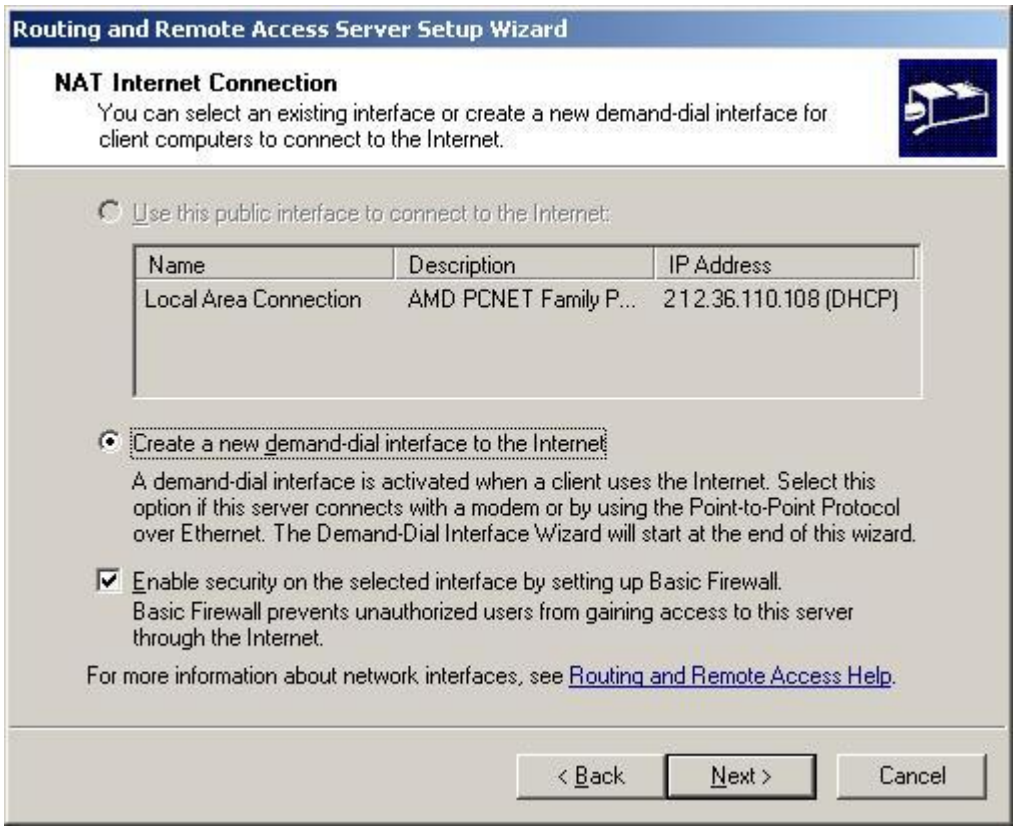

Press next to continue. The installation process will commence and services will be restarted, after which the finish screen will be displayed - showing what actions have taken place.

Configuring NAT

Configuration of NAT takes place from the Routing and Remote Access mmc found in the Administrative Tools folder in the Control Panel or on the start menu.

The screenshot below shows the routing and remote access mmc.

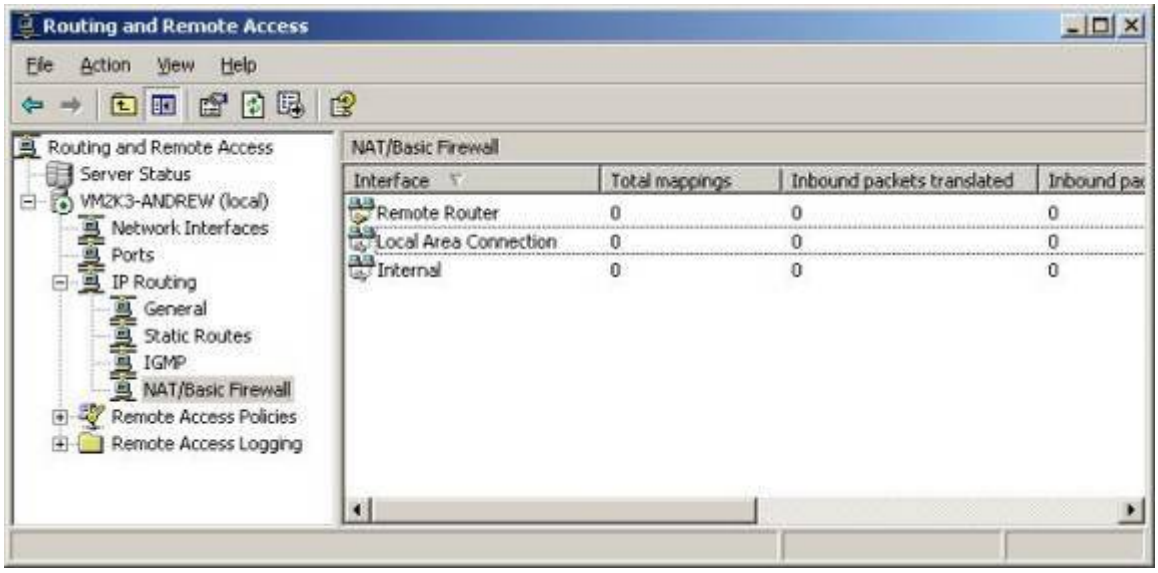

Select which interface you wish to configure and double click it. This will bring up the properties window giving you the option to change settings such as packet filtering and port blocking, as well as enabling/disabling certain features, such as the firewall.

The remote router (set up previously) properties box is shown below. The NAT/Basic Firewall tab is selected.

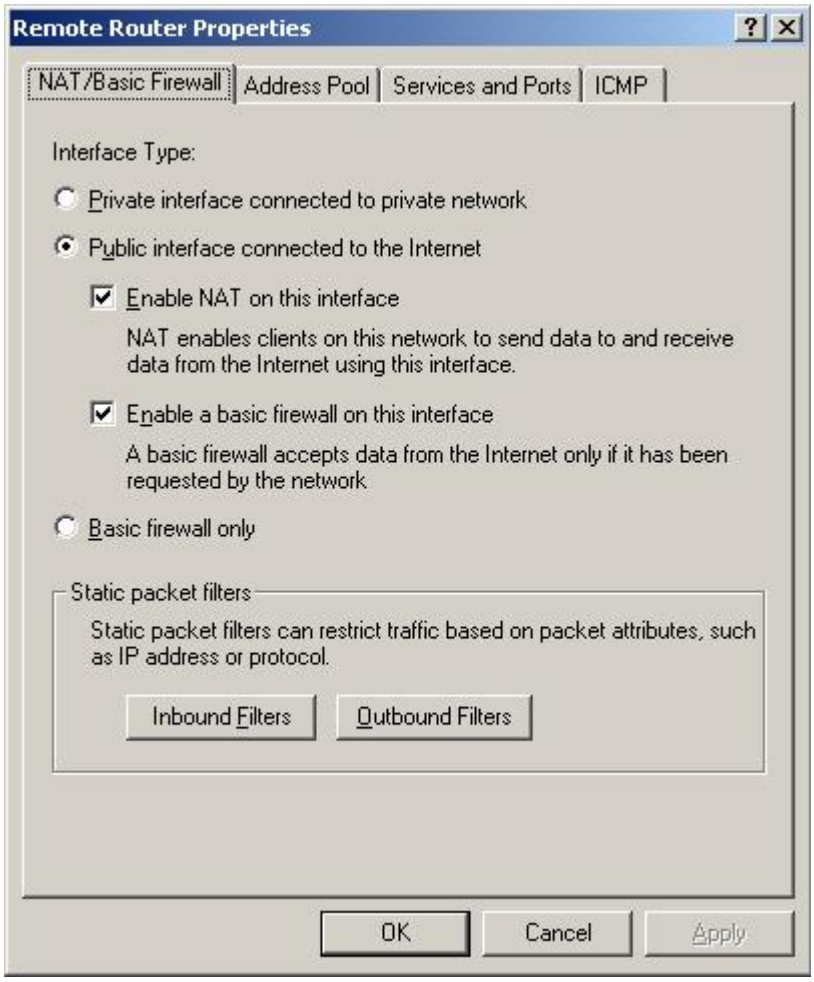

You are able to select the interface type – to specify what the network connection will be. In my example I have selected for the interface to be a public interface connected to the internet. NAT and the basic firewall option have also been enabled. The inbound and outbound buttons will open a window that will allow you restrict traffic based on IP address or protocol packet attributes. As per your instructions, certain TCP packets will be dropped before they reach the client computer. Thus, making the network safer and giving you more functionality. This is useful if, for example, you wanted to reject all packets coming from a blacklisted IP address or restrict internal users access to port 21 (ftp).

For further firewall configuration, go to the Services and Ports tab. Here you can select which services you would like to provide your users access to. You can also add more services by specifying details such as the incoming and outgoing port number.

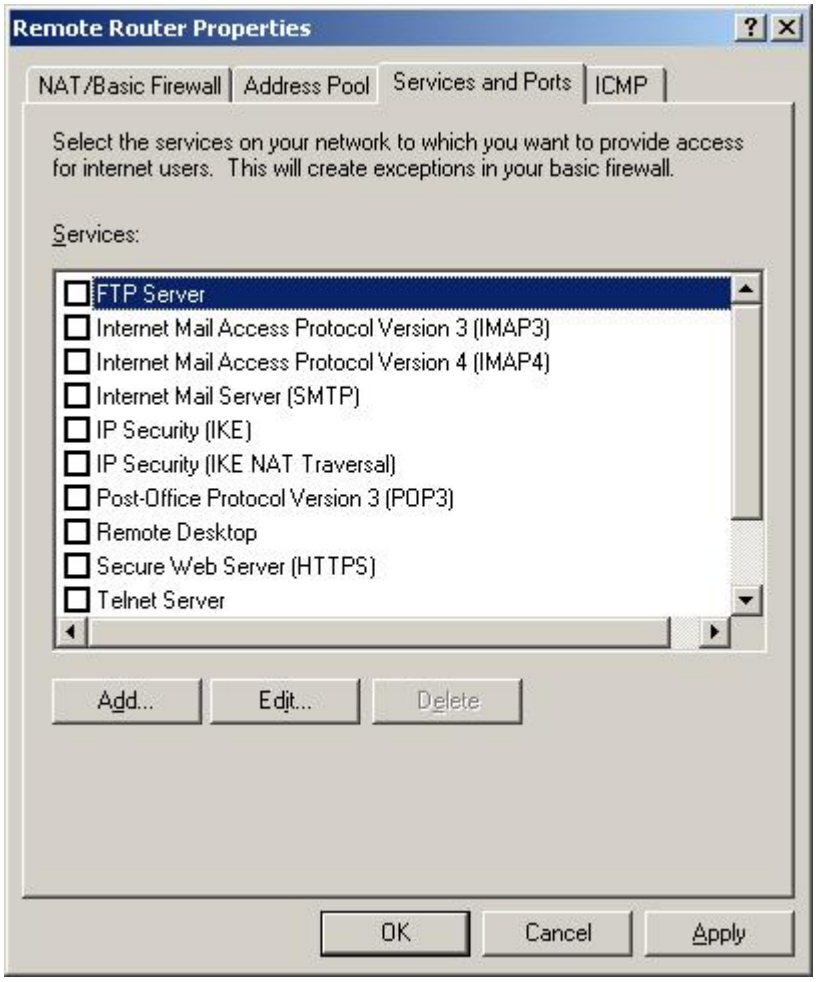

The list of services shown in the above screenshot are preset. Press Add to bring up the window that will allow the creation of a new service or select an available service and press Edit to modify that service. You will be asked to specify the name, TCP and UDP port number and the IP address of the computer hosting that service.

If the services in the list aren't enabled then any client computer on the Windows 2003 domain will not be able to access that specific service. For example, if the computer was configured as shown in the image above and a client computer tried to connect to an ftp site, he would be refused access. This section can prove to be very useful for any sized networks, but especially small ones.

That concludes this article. As you have seen, Network Address Translation is a useful feature that adds diversity and security to a network in a small to medium sized company. With the advent,

What is the real difference between NAT and PAT?
Port Address Translation (PAT) is a special kind of Network Address Translation (NAT). It can provide an excellent solution for a company that has multiple systems that need to access the Internet but that has only a few public IP addresses. Let's take a look at the distinctions between NAT and PAT and see how they are typically used. Then, I'll show you how to configure PAT on a Cisco router.

Understanding PAT and NAT

Before discussing PAT, it will help to describe what NAT does in general. NAT was designed to be a solution to the lack of public IP addresses available on the Internet. The basic concept of NAT is that it allows inside/internal hosts to use the private address spaces (10/8, 172.16/12, and 192.168/16 networks—see RFC1918), go through the internal interface of a router running NAT, and then have the internal addresses translated to the router's public IP address on the external interface that connects to the Internet.

If you dig into NAT a little deeper, you will discover that there are really three ways to configure it. From these configurations, you can perform a variety of functions. The three configurations are:

How do you allow inbound traffic for specific hosts on Windows 2003 NAT?

What is VPN? What types of VPN does Windows 2000 and beyond work with natively?

A. Microsoft defines a virtual private network as the extension of a private network that encompasses links across shared or public networks like the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link (such as a dial-up or long haul T-Carrier-based WAN link). Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated and encrypted is a VPN connection.

There are two key VPN scenarios—remote access and site-to-site. In remote access, the communications are encrypted between a remote computer (the VPN client) and the remote access VPN gateway (the VPN server) to which it connects. In site-to-site (also known as router-to-router), the communications are encrypted between two routers (VPN gateways) that link two sites.

What are the benefits of using VPN connections?

A. For remote access connections, an organization can use VPN connections to leverage the worldwide connectivity of the Internet and trade their direct-dial remote access solutions (and their corresponding equipment and maintenance costs) for a single connection to an Internet service provider (ISP) without sacrificing the privacy of a dedicated dial-up connection.

For routed connections, an organization can use VPN connections to leverage the worldwide connectivity of the Internet and trade long-distance dial-up or leased lines for simple connections to an Internet service provider (ISP) without sacrificing the privacy of a dial-up or dedicated site-to-site link.

What is IAS? In what scenarios do we use it?

A. IAS is the Windows implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy in Windows Server 2003.

In Windows Server 2008, the RADIUS server and proxy implementation is known as Network Policy Server (NPS).
What is IAS?

A. IAS is the Windows implementation of a Remote Authentication Dial-In User Service (RADIUS) server and proxy in Windows Server 2003.

In Windows Server 2008, the RADIUS server and proxy implementation is known as Network Policy Server (NPS).
Internet Authentication Service
Internet Authentication Service (IAS) in Microsoft® Windows Server® 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server and proxy. As a RADIUS server, IAS performs centralized connection authentication, authorization, and accounting for many types of network access including wireless, authenticating switch, and remote access dial-up and virtual private network (VPN) connections. As a RADIUS proxy, IAS forwards authentication and accounting messages to other RADIUS servers. RADIUS is an Internet Engineering Task Force (IETF) standard. For more detailed information, see Features of IAS
To optimize IAS authentication and authorization response times and minimize network traffic, install IAS on a domain controller.

When universal principal names (UPNs) or Windows Server 2003 domains are used, IAS uses the global catalog to authenticate users. To minimize the time it takes to do this, install IAS on either a global catalog server or a server that is on the same subnet. For more information, see The role of the global catalog. For more information about domain functionality, see Domain and forest functionality.

When you have remote RADIUS server groups configured and, in IAS Connection Request Policies, you clear the Record accounting information on the servers in the following remote RADIUS server group check box, these groups are still sent network access server (NAS) start and stop notification messages. This creates unnecessary network traffic. To eliminate this traffic, disable NAS notification forwarding for individual servers in each remote RADIUS server group by clearing the Forward network start and stop notifications to this server check box. For more information, see Configure the authentication and accounting settings of a group member and Configure accounting

What's the difference between Mixed mode and Native mode in AD when dealing with RRAS?

Like Windows 2000 and Active Directory, Exchange 2000 also has native and mixed modes of operation. Moving your Exchange organization to native mode offers advantages over mixed mode, but you must thoroughly understand the differences between native and mixed mode before planning a switch to native mode.

By default, Exchange 2000 installs and operates in mixed mode. Mixed mode allows Exchange 2000 and Exchange 5.5 servers to coexist and communicate. However, this backward compatibility limits administrative flexibility. Under mixed mode, Exchange 5.5 sites map directly to administrative groups and administrative groups map directly to Exchange 5.5 sites. All servers in a site must use a common service account, just as with Exchange 5.5. In addition, routing groups only contain servers from a single administrative group.

Native mode allows more flexibility than mixed mode. With Exchange in native mode, you can place servers from multiple administrative groups into a single routing group, and you can move servers between routing groups. You can do away with the requirement that all servers in a site must use a common service account. Additionally, operating in native mode allows you to move mailboxes between servers in the organization (removing the intersite mailbox move limitation in Exchange 5.5). For some companies, this enhanced mailbox move capability is reason enough to switch to native mode.

What's the difference between Mixed mode and Native mode in AD when dealing with RRAS?

The domain functional levels that can be set for Active Directory in Windows Server 2003 are listed below. The Windows 2000 Mixed and Windows Native domain functional levels were available in Windows 2000 to enable backward compatibility to operating systems such as Windows NT 4.0. The latter two functional levels are only available with Windows Server 2003.

- *Windows 2000 Mixed*: This is the default functional level implemented when you install a Windows Server 2003 domain controller. The basic Active Directory features are available when this mode is configured.
- *Windows 2000 Native*: In Windows 2000 Native functional level, the backup domain controllers of Windows NT is not supported as domain controllers in the domain. Only Windows 2000 domain controllers and Windows Server 2003 domain controllers are supported.

The main differences between Windows 2000 Mixed and Windows 2000 Native when discussing Active Directory features is that features like group nesting, or using Universal Groups and Security ID Histories (SIDHistory) is not available in Windows 2000 Mixed, but is available in Windows 2000 Native.

- *Windows Server 2003 Interim*: This functional level is used when Windows NT domains are directly upgraded to Windows Server 2003. Windows Server 2003 Interim is basically identical to Windows 2000 Native. The key point to remember on Windows Server 2003 Interim is that this functional level is used when the forests in your environment do not have Windows 2000 domain controllers.
- *Windows Server 2003*: This domain functional level is used when the domain only includes Windows Server 2003 domain controllers.

The features available for the new Windows Server 2003 Interim and Windows Server 2003 domain functional levels are discussed later on in this article.

The forest functional level can also be raised to enable additional Active Directory features. You have to though first raise the functional of domains within a forest before you can raise the forest functional level to Windows Server 2003. The domain functional level in this case has to be Windows 2000 Native or Windows Server 2003 before you raise the forest functional level. Domain controllers in the domains of the forest automatically have their functional level set to Windows Server 2003 when you raise the forest functional level to Windows Server 2003. Additional Active Directory features are immediately available for each domain in the forest.

The forest functional levels that can be set for Active Directory in Windows Server 2003 listed below.

- *Windows 2000*: In this forest functional level, Windows NT, Windows 2000 and Windows Server 2003 domain controllers can exist in domains.
- *Windows Server 2003 Interim*: Windows NT backup domain controllers and Windows Server 2003 domain controllers can exist in domains.
- *Windows Server 2003*: The domain controllers are all running Windows Server 2003.

Your Exchange organization is a candidate for native mode operation if you have no remaining Exchange 5.5 servers--or plans to add any--and you don't require Exchange 5.5 connectors.

Now that you know about native vs. mixed mode, you may want to start planning a switch to native mode. While making the switch isn't difficult, it's permanent. Begin testing and refining your plan for switching to native mode in a lab environment now.

What is the "RAS and IAS" group in AD?

The **Remote Access Service** (RAS) allows computers to make network connections to each other using telephone lines.

Windows 3.11, Windows 95, and Windows NT Workstation and Server can call out to another computer or communications server. Windows NT Workstation and Server can accept calls from any of these OS types. In most cases, the type of connection between the computers will be the Point-to-Point Protocol (PPP) running over an analog phone line, but RAS also supports the Serial Line Internet Protocol (SLIP), X.25, and ISDN connections.

This chapter describes the network protocols used by RAS, and how to install, configure, and start the RAS services. The syntax of RAS scripting languages is

RAS Protocols

A RAS connection running PPP can support the TCP/IP, IPX/SPX, and NetBEUI protocols simultaneously. For example, you may browse a remote network and share files using NetBEUI, browse the World Wide Web using TCP/IP, and access a Novell NetWare file server using IPX--all over a single dial-up RAS connection.

The SLIP protocol is TCP/IP-only and does not support the multitude of options available in PPP. It has fallen out of favor for these reasons and can be avoided in most situations.

The NetBEUI protocol will not generally be able to run over the RAS connection unless you have a Microsoft operating system at both ends of the connection. NetBIOS functions such as file sharing, printing, and NetBIOS name service can all be made to work over either a TCP/IP-only PPP or a SLIP connection using NetBIOS over TCP/IP (NBT).

What are Conditions and Profile in RRAS Policies?

What is remote access policy? What is its usage?
Remote access policies are an ordered set of rules that define whether remote access connection attempts are either authorized or rejected. Each rule includes one or more **conditions** (which identifies the criteria), a set of **profile** settings (to be applied on the connection attempt), and a **permission** setting (grant or deny) for remote access.

This can be compared like a brain of the door-keeper (VPN server) which allows entry to your network from outside . Remote access policy decides who can access what resources from where using what tunnel settings. So configuring proper set of policies are important.

What are some common examples?
You may want to have different policies based upon one or more factors in different **conditions**:
1) Who is accessing the network(Windows-Groups)
2) What tunnel type is getting used (Tunnel-Type)
3) What authentication type is getting used
4) What is the client's IP address (useful for site-to-site scenarios where the IP address of calling router remains same)
5) What time of day client is accessing (like you may want to block access at particular times)
6) etc

You may want to enforce following profile on a given policy:
1) Idle time after which the connection should be disconnected
2) Session time after which the connection should be disconnected
3) Inbound/Outbound filters that can be applied per PPP connection (or per user connection) - to restrict access of a given client/site to a given network (IP address, port number)
4) Encryption Type
5) Authentication Algorith Type

What types or authentication can a Windows 2003 based RRAS work with?

The previous article in this series discussed improvements made to IPSec in Windows Server 2003. This article expands on this topic with an examination of the security-related enhancements made in Routing and Remote Access Service (RRAS) and Internet Authentication Service (IAS), specifically:

- Support for L2TP/IPSec over NAT
- Network Access Quarantine
- NetBIOS-related enhancements
- EAP-TLS improvements
- Improved remote access client support
- IAS Proxy

RRAS was introduced as a built-in component in Windows 2000 Server (but it is also available as an add-on for Windows NT 4.0 Server). As its name indicates, it combines routing and remote access functionality into a single administrative interface, allowing the server to be turned into a secure, software-based router or a remote access server, or both. IAS (which first appeared in Windows 2000 server) is Microsoft's implementation of Remote Authentication Dial-In User Service (RADIUS), and its primary purpose is to provide authentication, authorization, and accounting functionality for remote access. Because of its role, it closely interacts with RRAS. Hence, this article describes both.

Windows 2003 RRAS has a number of new, nonsecurity-related features. It supports Point-to-Point Protocol over Ethernet (PPPoE), reflecting the growing popularity of broadband communication. It can also function as a bridge, combining separate, mixed media segments into a single networking subnet. What might also be a bit of surprise is the dependency

between RRAS and Internet Connection Firewall (ICF), since this component was not available in Windows 2000.

Like its Windows XP equivalent, the new version of ICF operates as a stateful firewall (intended for protecting Internet Connection Sharing), which means it tracks sessions initiated from the internal network and, by default, permits inbound traffic only if it constitutes part of these sessions. In addition, ICF selectively permits incoming traffic based on the targeted port and redirects it to any of internal IP addresses (on the same or a different port). Since the same functionality can be provided by RRAS (configurable from the NAT/Basic firewall tab of the interface properties dialog box in IP Routing node of the Routing and Remote Access MMC console snap-in), Microsoft decided to make them mutually exclusive. However, ICF must be disabled to activate RRAS to take full advantage of the security-related features detailed below.

How does SSL work?

SSL stands for " Secure Sockets layer " Socute is a technical term that referes to an application programming interface or API which that refers to an application prograaming interface or API which is used to communicate b/w to computer Layer refers to the level or layer of this communication b/w the computer.

Though it is good to answer the "How does SSL work?" question (see the steps on the following pages) the typical merchant really needs to only be concerned with how to get a secure certificate and making sure that he/she is using a valid and current ssl certificate (step 2.03) and what URL to use when creating secure links. SSL certificates are purchased from various certificate vendors and it requires a CSR (Certificate Signing Request) to be generated on the web server. This usually involves getting in touch with the hosting company and asking them to generate the CSR for you. After your receive the CSR (which looks like an encrypted block of undecipherable text) you can order your certificate from the SSL certificate provider. Once you receive the SSL certifcate back from the certificate authority, you will normally need the hosting company to install it for you.

How does IPSec work?

IPSec is an Internet Engineering Task Force (IETF) standard suite of protocols that provides data authentication, integrity, and confidentiality as data is transferred between communication points across IP networks. IPSec provides data security at the IP packet level. A packet is a data bundle that is organized for transmission across a network, and it includes a header and payload (the data in the packet). IPSec emerged as a viable network security standard because enterprises wanted to ensure that data could be securely transmitted over the Internet. IPSec protects against possible security exposures by protecting data while in transit.

How do I deploy IPSec for a large number of computers?

What is IPSec?

IPSec is an encryption protocol designed to work at the IP level. As you might know, Kerberos is the primary Windows authentication protocol. Kerberos and IPSec differ in that Kerberos provides user-to-service authentication. IPSec on the other hand is used to encrypt and authenticate communications between computers on the network. It is a low-level protocol that has absolutely nothing to do with securing access to data or services on a server.

IPSec's main goals are to encrypt communications across an IP-based network (such as the Internet and most private networks) and to guarantee that a transmission has not been tampered with en route.

One simple benefit of IPSec is that it's built into Windows. That means you don't have to buy any additional software and you don't have to worry about compatibility issues when implementing IPSec policies. You also don't have to do anything to deploy IPSec onto the server or client PCs -- you just create an appropriate group policy.

4. **Set policies for any Microsoft version**
Microsoft originally released IPSec with Windows 2000. This means that Windows 2000 (Server and Professional), Windows XP and Windows Server 2003 all support IPSec, but Windows 9.x does not. Fortunately, enabling IPSec does not require you to alienate Windows 9.x machines or machines running other operating systems that may not support IPSec. When you create the IPSec group policy entry, you can choose to have machines request security or require security.

If an IPSec policy is set to request security, a client that tries to communicate with the server will receive a request from that server to use IPSec communications. If the client supports IPSec, encrypted communications begin. If the client does not support IPSec, communications remain unencrypted. But if the IPSec policy requires security, all conversations must be encrypted by IPSec.

Generally speaking, setting up a security policy that requests IPSec security is perfect for most companies because it accomodates both IPSec-aware and non-IPSec-aware clients. As legacy operating systems are phased out, the newer operating systems will already be prepared to have secure communications with other machines.

What types of authentication can IPSec use?

Extended Authentication (XAUTH).

Extended Authentication (XAUTH) and Mode Configuration (MODE-CFG)

Authentication schemes such as Remote Authentication Dial-In User Service (RADIUS) and SecureID are commonly used for providing secure remote access. It is highly desirable to leverage these authentication mechanisms for IPSec remote access. But Internet Key Exchange (IKE) protocol, which you learned about in Chapter 2, "IPSec Overview," does not provide a method to leverage these unidirectional authentication schemes. Extended Authentication, commonly referred to as XAUTH, was developed to leverage these legacy authentication schemes with IKE.

XAUTH provides an additional level of authentication by allowing the IPSec gateway to request extended authentication from remote users, thus forcing remote users to respond with their credentials before being allowed access to the VPN. It should be noted that XAUTH functions by first forming an IKE phase 1 SA using conventional IKE, and then by extending the IKE exchange to include additional user authentication exchanges. Figure 4-1 shows an XAUTH exchange using a generic username and password authentication scheme

What is PFS (Perfect Forward Secrecy) in IPSec?

PFS

With PFS disabled, initial keying material is "created" during the key exchange in phase-1 of the IKE negotiation. In phase-2 of the IKE negotiation, encryption and authentication session keys will be extracted from this initial keying material. By using PFS, Perfect Forwarding Secrecy, completely new keying material will always be created upon re-key. Should one key be compromised, no other key can be derived using that information.

PFS can be used in two modes, the first is PFS on keys, where a new key exchange will be performed in every phase-2 negotiation. The other type is PFS on identities, where the identities are also protected, by deleting the phase-1 SA every time a phase-2 negotiation has been finished, making sure no more than one phase-2 negotiation is encrypted using the same key.

PFS is generally not needed, since it is very unlikely that any encryption or authentication keys will be compromised.

How do I monitor IPSec?

The IP Security Monitor snap-in, a new feature in Windows Server 2003, can be used to monitor and troubleshoot IPSec activity. The IP Security Monitor snap-in provides enhanced IPSec security monitoring. As long as the IPSec policy is active, you can monitor how the IPSec policy is functioning within your networking environment through the IP Security Monitor.

The main administrative activities which you can perform through the IP Security Monitor snap-in are listed here:

- Customize the IP Security Monitor display
- Monitor IPSec information on the local computer.
- Monitor IPSec information on remote computers.
- View IPSec statistics.
- View information on IPSec policies
- View security associations information.
- View generic filters
- View specific filters
- Search for specific filters based on IP address

By default, the computer which is listed in the IP Security Monitor snap-in is the local computer. You can though add another computer(s) which you want to monitor to the IP Security Monitor

Looking at IPSec-encrypted traffic with a sniffer. What packet types do I see?

Mirrored Packet u can see

What can you do with NETSH?

What can we do with Netsh.exe?

With Netsh.exe you can view your TCP/IP settings.

Type the following command in a Command Prompt window (CMD.EXE):

netsh interface ip show config

You can configure your computer's IP address and other TCP/IP related settings. For example:

The following command configures the interface named Local Area Connection with the static IP address 192.168.0.100, the subnet mask of 255.255.255.0, and a default gateway of 192.168.0.1:

netsh interface ip set address name="Local Area Connection" static 192.168.0.100 255.255.255.0 192.168.0.1 1

(The above line is one long line, watch for word wrap. Copy paste it as one line)

Netsh.exe can be useful in certain situations when you have a portable computer that needs to be relocated between 2 or more office locations, while still maintaining a specific and static IP address configuration. With Netsh.exe, you can save and restore the appropriate network configuration all from the command prompt.

Connect your portable computer to location #1, and then manually configure the required network settings.

Now, you need to export your current IP settings to a text file. Use the following command:

netsh -c interface dump > c:\location1.txt

When you reach location #2, do the same thing, only keep the new settings to a different file:

netsh -c interface dump > c:\location2.txt

You can go on with as many other location you may need.

Now, whenever you need to travel between locations, you can enter the following command in a Command Prompt window (CMD.EXE):

netsh -f c:\location1.txt or **netsh -f c:\location2.txt**

Netsh.exe can also be used to configure your NIC to automatically obtain an IP address from a DHCP server:

netsh interface ip set address "Local Area Connection" dhcp

You can use this command to setup WINS:

netsh interface ip set wins "Local Area Connection" static 192.168.0.200

Or, if you want, you can configure your NIC to dynamically obtain it's DNS settings:

netsh interface ip set dns "Local Area Connection" dhcp

Netsh is very customizable and very useful.

How do I look at the open ports on my machine?

Shell to "Netstat" and save the result as a string and rip it line by line

Code:

```
Call Shell("command.com /c netstat -an -o > " & sfile, vbNormal)
```

2. Use GetTcpTable() and GetUdpTable(), as mike said or you can use
the undocumented AllocateAndGetTCPExTableFromStack API to get
the same list but with the PID (work only with XP).

3. Use the Native API, not documented, NtQuerySystemInformation().
To know the processes with open ports.
With this API you will access the TDI level (Transport Driver Interface) located in the system library NTDLL.DLL

Technical Interview Questions – Active Directory

What is domain ?

Answer

In Windows NT and Windows 2000, a domain is a set of network resources (applications, printers, and so forth) for a group of users. The user need only to log in to the domain to gain access to the resources, which may be located on a number of different servers in the network.

The 'domain' is simply your computer address not to confused with an URL. A domain address might look something like 211.170.469.

What is domain controller ?

Primary domain controller (PDC) and backup domain controller (BDC) are roles that can be assigned to a server in a network of computers that use the Windows NT operating system. Windows NT uses the idea of a domain to manage access to a set of network resources (applications, printers, and so forth) for a group of users. The user need only to log in to the domain to gain access to the resources, which may be located on a number of different servers in the network. One server, known as the primary domain controller, manages the master user database for the domain. One or more other servers are designated as backup domain controllers. The primary domain controller periodically sends copies of the database to the backup domain controllers. A backup domain controller can step in as primary domain controller if the PDC server fails and can also help balance the workload if the network is busy enough.

Setting up and maintaining PDCs and BDCs and domain information is a major activity for the administrator of a Windows NT network. In Windows 2000, the domain controller concept is retained but the PDC and BDC server roles are generally replaced by the *Active Directory*.

What is domain tree ?

Domain Trees

A *domain tree* comprises several domains that share a common schema and configuration, forming a contiguous namespace. Domains in a tree are also linked together by trust relationships. Active Directory is a set of one or more trees.

Trees can be viewed two ways. One view is the trust relationships between domains. The other view is the namespace of the domain tree.

What is forests ?

A collection of one or more domain trees with a common schema and implicit trust relationships between them. This arrangement would be used if you have multiple root DNS addresses.

An Introduction to the Active Directory Features

With the release of Microsoft Windows Server 2003 quite a few enhancements and features were introduced that were not previously available in Windows 2000. These enhancements were aimed at improving the scalability, efficiency, speed and performance of Active Directory, and addressed a few deficiencies or shortcomings of the earlier version of Active Directory utilized in Windows 2000 Server.

When a domain controller running Windows Server 2003 is created, a number of Active Directory basic features are immediately installed and available to the Windows Server 2003 domain controller. Certain other Active Directory features are only available when particular conditions exist in the network.

Additional Active Directory features can be enabled but is dependant on the following conditions, or factors:

- The operating system (OS) running on the domain controller
- The domain functional level. In Windows 2000 Active Directory, the domain mode terminology was utilized.
- The forest functional level
- Whether the functional level is raised for the domain only, or for the forest.

The domain functional levels that can be set for Active Directory in Windows Server 2003 are listed below. The Windows 2000 Mixed and Windows Native domain functional levels were available in Windows 2000 to enable backward compatibility to operating systems such as Windows NT 4.0. The latter two functional levels are only available with Windows Server 2003.

- *Windows 2000 Mixed*: This is the default functional level implemented when you install a Windows Server 2003 domain controller. The basic Active Directory features are available when this mode is configured.
- *Windows 2000 Native*: In Windows 2000 Native functional level, the backup domain controllers of Windows NT is not supported as domain controllers in the domain. Only Windows 2000 domain controllers and Windows Server 2003 domain controllers are supported.

The main differences between Windows 2000 Mixed and Windows 2000 Native when discussing Active Directory features is that features like group nesting, or using Universal Groups and Security ID Histories (SIDHistory) is not available in Windows 2000 Mixed, but is available in Windows 2000 Native.

- *Windows Server 2003 Interim*: This functional level is used when Windows NT domains are directly upgraded to Windows Server 2003. Windows Server 2003 Interim is basically identical to Windows 2000 Native. The key point to remember on Windows Server 2003 Interim is that this functional level is used when the forests in your environment do not have Windows 2000 domain controllers.
- *Windows Server 2003*: This domain functional level is used when the domain only includes Windows Server 2003 domain controllers.

The features available for the new Windows Server 2003 Interim and Windows Server 2003 domain functional levels are discussed later on in this article.

The forest functional level can also be raised to enable additional Active Directory features. You have to though first raise the functional of domains within a forest before you can raise the forest functional level to Windows Server 2003. The domain functional level in this case has to be Windows 2000 Native or Windows Server 2003 before you raise the forest functional level. Domain controllers in the domains of the forest automatically have their functional level set to Windows Server 2003 when you raise the forest functional level to Windows Server 2003. Additional Active Directory features are immediately available for each domain in the forest.

The forest functional levels that can be set for Active Directory in Windows Server 2003 listed below.

- *Windows 2000*: In this forest functional level, Windows NT, Windows 2000 and Windows Server 2003 domain controllers can exist in domains.
- *Windows Server 2003 Interim*: Windows NT backup domain controllers and Windows Server 2003 domain controllers can exist in domains.
- *Windows Server 2003*: The domain controllers are all running Windows Server 2003.

New Active Directory Basic Features

Active Directory basic features are enabled by default when you install a Windows Server 2003 domain controller. The enhancements and features available are summarized below:

- You can promote domain controllers to Windows Server 2003 domains more efficiently and faster because you can use a tape backup of the Active Directory database which is essentially a restored backup of another domain controller, to update the Active Directory database for a newly promoted domain controller. This *decreases the time needed to install an additional domain controller in an existing domain*.
- Because problems do at times presents themselves when Windows NT 4 primary domain controllers are upgraded to Windows Server 2003 domain controllers, you can configure the domain controllers to treat your Windows clients as Windows NT domain controllers.
- Active Directory can now store over one billion objects, thereby *improving scalability*.
- In Active Directory used in Windows 2000, changes made to the identical Group hosted on multiple domain controllers in the same replication interval used to overwrite each other. This has since been corrected as *group members are replicated as separate entities.*
- The actual *method used to calculate the replication topology between sites is streamlined* to solve a prior problem whereby the replication topology calculations could not be completed in the specified time.
- *Domain logon has been improved*, and users can continue to log on at times when the Global Catalog server cannot be accessed because Universal group membership can now be stored on servers that are not Global Catalog servers.
- Windows Server 2003 introduces a new naming context, or directory partition, namely the *Application directory partition.* Application specific data is stored in this directory partition. You can now configure replication for application specific data between domain controllers. The Application directory partition is primarily used to store DNS record objects for Active Directory Integrated zones.
- The *inetOrgPerson object class is a new security principal* added to the base schema. You use this security principal in the same manner that you would use other security principals such as User and Group.
- With Windows Server 2003, support is included for:
    - RFC 2589 - LDAPv3: Extensions for Dynamic Directory Services Two private addresses are utilized for communication among the nodes: You can now store information that is time sensitive in Active Directory.
    - RFC 2829 - Authentication Methods for LDAP: It is now simpler to integrate Active Directory into environments that are not running Windows.
    - RFC 2830 - LDAPv3: Extension for Transport Layer Security: Secure connections are now used when Lightweight Directory Access Protocol (LDAP) queries are transmitted over the network to domain controllers. Active Directory encrypts all LDAP traffic by default.

- *Enhancements to LDAP queries* include a new query types called an *Attribute Scoped Query (ASQ)*; and a new LDAP management mechanism called *Virtual List Views*. You can use the ASQ to determine those groups to which a particular user is a member of. Virtual List Views enable you to view a large set of data in an order.
- Active Directory quotas can be used to control and manage the number of objects that a user, group, and computer can be the owner of in a particular Active Directory directory partition.
- Because you now simultaneously select multiple directory objects, you can *simultaneously change the attributes on multiple objects*.
- You can also use the *new drag-and-drop move feature* to move directory objects from one container to another container. You can use the same feature to add objects to group membership lists.
- With Windows Server 2003, *you can save, export and refresh Active Directory queries*. Through the use of saved queries, you can find specific objects, and modify the properties of these objects simultaneously.
- You can use the following *Windows Server 2003 Active Directory command-line tools* to administer Active Directory:
    - Dsadd, Dsget, Dsmod, Dsmove, Dsquery, Dsrm, Csvde, Ntdsutil, Ldifde
- A *new version of the Active Directory Migration Tool (ADMT)* includes the following:
    - Password migration support
    - Access to user profiles remain unchanged
- With the introduction of Windows Server 2003 Active Directory came the introduction of more than 200 *new Group Policy settings*. The new Group Policy settings have to though be applied to Windows Server 2003 clients for it to be enabled.
- Active Directory in Windows Server 2003 has an *integrated Resultant Set of Policies (RSoP) calculator* that can be used to determine the policies which have been applied to a particular user or computer. You can use the feature through the Resultant Set Of Policy (RSoP) Wizard or from the command-line.
- Windows Server 2003 Help and the Group Policy console (Extended tab) now include descriptive information on all the administrative templates.
- The following *new command-line tools can be used to manage Group Policy*:
    - Gpupdate: The Gpupdate update tool replaces the Secedit switches used in Windows 2000. You can use Gpupdate to immediately refresh group policy.
    - Gpresult: The Gpresult tool is used to create and view the results of a RSoP query using command line.
- When deploying software with Group Policy, you can force assigned applications to be installed at deployment, and you can now choose to enable or disable the following advanced options:
    - The publication of OLE class information on a software package
    - The availability of 32-bit programs to 64-bit computers

Active Directory Features enabled by Domain/Forest Functional Levels

The features and enhancements listed below are only available when the domain or forest functional levels have been raised to Windows Server 2003. This means that each domain controller should be running Windows Server 2003. The features and enhancements enabled by this functional level can be used to change the configuration of the domain and forest.

- *Domain controller renaming tool*: You can use the domain controller renaming tool to rename domain controllers – you do not need to first demote them. All Active Directory and DNS entries are automatically updated as well.

- *Domain rename utility (Rendom.exe)*: You can use the Rendom.exe utility to change the name of domains. Through the utility, you can change the NetBIOS or DNS name of a domain. This includes any child, parent, domain tree root, or forest root domain.
- You can now *restructure your forest* by moving existing domains to different locations in the domain hierarchy.
- With the functional level raised to Windows Server 2003, you can create forest trust to form a *two-way transitive trust relationship between two forests*. This trust relationship enables users in one forest to access resources available in the forest.
- For the Active Directory schema, enhancements include the capability of *now assigning an auxiliary schema class to a specific object(s)*. The support feature is called dynamically linked auxiliary classes.
- When Active Directory schema objects are no longer needed, you can *disable classes and attributes, rename classes and attributes and redefine them*. You can also re-activate these classes and attributes when you need them at a later date. You cannot however delete schema objects.
- You can also restrict users in a particular domain or forest from accessing network resources in a different domain/forest. By controlling resource access between domains and forests, you can allow users specific access to network resources.
- Global catalog replication has also been improved. When there is an extension of the partial attribute set, only the attributes which have been added, are replicated. This in turn decreases the amount of traffic generated by global catalog replication.

What is LDAP?

LDAP (Lightweight Directory Access Protocol) is a protocol for communications between LDAP servers and LDAP clients. LDAP servers store "directories" which are access by LDAP clients.

LDAP is called *lightweight* because it is a smaller and easier protocol which was derived from the X.500 DAP (Directory Access Protocol) defined in the OSI network protocol stack.

LDAP servers store a hierarchical directory of information. In LDAP parlance, a fully-qualified name for a directory entry is called a *Distinguished Name*. Unlike DNS (Domain Name Service) FQDN's (Fully Qualified Domain Names), LDAP DN's store the most significant data to the right.

The Four Models of LDAP

LDAP is defined by four models:

Model	Description
Information	Describes the structure of information stored in an LDAP directory
Naming	Describes how information in an LDAP directory is organized and identified
Functional	Describes what operations can be performed on the information stored in an LDAP directory
Security	Describes how the information in an LDAP directory can be protected from unauthorized access

LDAP is extensible and can be used to store any type of data. Most interesting is that LDAP is being used as a core technology for most Single Sign On (SSO) implementations.

Can you connect Active Directory to other 3rd-party Directory Services? Name a few options.

Yes, you can use dirXML or LDAP to connect to other directories (ie. E-directory from Novell).

Where is the AD database held? What other folders are related to AD?

Active Directory Structure

Active Directory has a hierarchical structure that consists of various components which mirror the network of the organization. The components included in the Active Directory hierarchical structure are listed below:

- Sites
- Domains
- Domain Trees
- Forests
- Organizational Units (OUs)
- Objects
- Domain Controllers
- Global Catalog
- Schema

The Global Catalog and Schema components actually manage the Active Directory hierarchical structure. In Active Directory, logically grouping resources to reflect the structure of the organization enables you to locate resources using the resource's name instead of its physical location. Active Directory logical structures also enable you to manage network accounts and shared resources.

The components of Active Directory that represent the *logical structure* in an organization are:

- Domains, Organizational Units (OUs), Trees, Forests, Objects

The components of Active Directory that are regarded as Active Directory physical structures are used to reflect the organization's physical structure. The components of Active Directory that are *physical structures* are:

- Sites, Subnets, Domain Controllers

The following section examines the logical and physical components of Active Directory.

A *domain* in Active Directory consists of a set of computers and resources that all share a common directory database which can store a multitude of objects. Domains contain all the objects that exist in the network. Each domain contains information on the objects that they contain. In Active Directory, domains are considered the core unit in its logical structure. Domains in Active Directory actually differ quite substantially from domains in Windows NT networks. In Windows NT networks, domains are able to store far less objects than what Active Directory domains can store. Windows NT domains are structured as peers to one another. What this means is that you cannot structure domains into a hierarchical structure. Active Directory domains on the other hand can be organized into a hierarchical structure through the use of forests and domain trees.

An Active Directory domain holds the following:

- Logical partition of users and groups
- All other objects in the environments

In Active Directory, domains have the following common characteristics:

- The domain contains all network objects
- The domain is a security boundary – access control lists (ACLs) control access to the objects within a domain.

Within a domain, objects all have the following common characteristics:

- Group Policy and security permissions
- Hierarchical object naming
- Hierarchical properties
- Trust relationships

The majority of components in Active Directory are *objects*. In Active Directory, objects represent network resources in the network. Objects in Active Directory have a unique name that identifies the object. This is known as the *distinguished name* of the object. Objects can be organized and divided into *object classes*. Object classes can be regarded as the logical grouping of objects. An object class contains a set of *object attributes* which are characteristics of objects in the directory. Attributes can be looked at as properties that contain information on characteristics and configurations. The Active Directory objects that an Administrator would most likely be concerned with managing are users, groups and computers. In Active Directory, the main groups are *security groups* and *distribution groups.* It is easier to place users into groups and then assign permissions to network resources via these groups. Through implementing groups and using groups effectively, you would be in a good position to manage security and permissions in Active Directory.

Organizational units (OUs) can be considered logical units that can be used to organize objects into logical groups. OUs can be hierarchically arranged within a domain. An organization unit can contain objects such as user accounts, groups, computers, shared resources, and other OUs. You can also assign permissions to OUs to delegate administrative control. Domains can have their own OU hierarchy. Organizational units are depicted as folders in the Active Directory Users And Computers administrative tool.

In Active Directory, a *domain tree* is the grouping of one or multiple Windows 2000 or Windows Server 2003 domains. Domain trees are essentially a hierarchical arrangement of these domains. Domain trees are created by adding child domains to a parent domain. Domains that are grouped into a domain tree have a hierarchical naming structure and also share a contiguous namespace.

Multiple domains are typically utilized to:

- Improve performance
- Decentralize administration
- Manage and control replication in Active Directory
- Through the utilization of multiple domains, you can implement different security policies for each domain.
- Multiple domains are also implemented when the number of objects in the directory is quite substantial.

A *forest* in Active Directory is the grouping of one or multiple domain trees. The characteristics of forests are summarized below:

- Domains in a forest share a common schema and global catalog, and are connected by implicit two-way transitive trusts. A global catalog is used to increase performance in Active Directory when users search for attributes of an object. The global catalog server contains a copy of all objects in its associated host domain, as well as a partial copy of objects in the other domains in the forest.
- Domains in a forest function independently, with the forest making communication possible with the whole organization.
- Domain trees in a forest do not have the same naming structures.

In Active Directory, a *site* is basically the grouping of one or more Internet Protocol (IP) subnets which are connected by a reliable high-speed link. Sites normally have the same boundaries as a local area network (LAN). Sites should be defined as locations that enable fast and cheap network access. Sites are essentially created to enable users to connect to a domain controller using the reliable high-speed link; and to optimize replication network traffic. Sites determine the time and the manner in which information should be replicated between domain controllers.

A site contains the objects listed below that are used to configure replication among sites.

- Computer objects
- Connection objects

A *domain controller* is a computer running Windows 2000 or Windows Server 2003 that contains a replica of the domain directory. Domain controllers in Active Directory maintain the Active Directory data store and security policy of the domain. Domain controllers therefore also provide security for the domain by authenticating user logon attempts. The main functions of domain controllers within Active Directory are summarized in the following section:

- Each domain controller in a domain stores and maintains a replica of the Active Directory data store for the particular domain.
- Domain controllers in Active Directory utilize multimaster replication. What this means is that no single domain controller is the master domain controller. All domain controllers are considered peers.
- Domain controllers also automatically replicate directory information for objects stored in the domain between one another.
- Updates that are considered important are replicated immediately to the remainder of the domain controllers within the domain.
- Implementing multiple domain controllers within a domain provides fault tolerance for the domain.
- In Active Directory, domain controllers can detect collisions. Collisions take place when an attribute modified on one particular domain, is changed on a different domain controller prior to the change on the initial domain controller being fully propagated.

Apart from domain controllers, you can have servers configured in your environment that operate as *member servers* of the domain but who do not host Active Directory information. Member servers do not provide any domain security functions either such as authenticating users. Typical examples of member servers are file servers, print servers, and Web servers.

Standalone severs on the other hand operate in workgroups and are not members of the Active Directory domain. Standalone servers have, and manage their own security databases.

Active Directory Namespace Structure

The Domain Name System (DNS) is the Internet service that Active Directory utilizes to structure computers into domains. DNS domains have a hierarchical structure that identifies computers, organizational domains and top-level domains. Because DNS also maps host names to numeric Transmission Control Protocol/Internet Protocol (TCP/IP) addresses, you define the Active Directory domain hierarchy on an Internet-wide basis, or privately. Because DNS is an important component of Active Directory, it has to be configured before you install Active Directory.

The information typically stored in Active Directory can be categorized as follows:

- *Network security entities*: This category contains information such as users, groups, computers, applications.
- *Active Directory mechanisms*: This category includes permissions, replication, and network services.
- *Active Directory schema*: Active Directory objects that define the attributes and classes in Active Directory are included here.

To ensure compatibility with the Windows NT domain model, Active Directory is designed and structured on the idea of domains and trust relationships. Because the SAM databases in Windows NT could not be combined, domains have to be joined using trust relationships.

With Active Directory, a domain defines the following:

- A namespace
- A naming context
- A security structure
- A management structure

Within the domain, you have users and computers that are members of the domain, and group policies. In Active Directory, you can only create a naming context at a domain boundary, or by creating an Application naming context. An *Application naming context* is a new Active Directory feature introduced in Windows Server 2003. Other than a *Domain naming context*, each installation of Active Directory must have a Schema naming context, and a Configuration naming context.

- *Schema naming context*: Domain controllers in the forest each have a read-only replica of the Schema naming context which contains the ClassSchema and AttributeSchema objects. These objects signify the classes and attributes in Active Directory. The domain controller acting the role of Schema Role Master is the only domain controller that can change the schema.
- *Configuration naming context*: Domain controllers in the forest each have a read and write replica of the Configuration naming context. The Configuration naming context contains the top-level containers listed below which basically manage those services that support Active Directory:
    - *Display Specifiers container*: Objects which change the attributes that can be viewed for the remainder of the object classes are stored in this container. Display Specifiers supply localization and define context menus and property pages. Localization deals with determining the country code utilized during installation, and then moves all content via the proper Display Specifier.

Context menus and property pages are defined for each user according to whether the user attempting to access a particular object has Administrator privileges.

- *Extended Rights containe*r: Because you can assign permissions to objects and the properties of an object, Extended Rights merges various property permissions to form a single unit. In this manner, Extended Rights manages and controls access to objects.
- *Lost and Found Config container*: The Domain naming context and Configuration context each have a Lost and Found Config container that holds objects which have gone astray.
- *Partitions container*: The Partitions container contains the cross-reference objects that depict all the other domains in a forest. The Partitions container's data is referenced by domain controllers when they create referrals to these domains. The data in the Partitions container can only be altered by a single domain controller within the forest.
- *Physical Locations container*: The Physical Locations container contains physical Location DN objects which are related to Directory Enabled Networking (DEN).
- *Services container*: This container stores the objects of distributed applications and is replicated to all domain controllers within the forest. You can view the contents of the container in the Active Directory Sites and Services console.
- *Sites container*: The objects stored in the Sites container control Active Directory replication, among other site functions. You can also view the contents of this container in the Active Directory Sites and Services console.
- *Well-Known Security Principals container*: This container stores the names and unique Security Identifiers (SIDs) for groups such as Interactive and Network.

Replication and Active Directory

In Active Directory, directory data that is classified into the categories listed below are replicated between domain controllers in the domain:

- *Domain data* includes information on the objects stored in a particular domain. This includes objects for user accounts, Group Policy, shared resources and OUs.
- *Configuration data* includes information on the components of Active Directory that illustrates the structure of the directory. Configuration data therefore define the domains, trees, forests and location of domain controllers and global catalog servers.
- *Schema data* lists the objects and types of data that can be stored in Active Directory.

Active Directory utilizes *multimaster replication*. This means that changes can be made to the directory from any domain controller because the domain controllers operate as peers. The domain controller then replicates the changes that were made. Domain data is replicated to each domain controller within that domain. Configuration data and schema data are replicated to each domain in a domain tree and forest. Objects stored in the domain are replicated to global catalogs. A subset of object properties in the forest is also replicated to global catalogs. Replication that occurs within a site is known as *intra-site replication*. Replication between sites is known as *inter-site replication*.

Support Files of Active Directory

The Active Directory support files are listed below. These are the files that you specify a location for when you promote a server to a domain controller:

- *Ntds.dit (NT Directory Services)*: Ntds.dit is the core Active Directory database. This file on a domain controller lists the naming contexts hosted by that particular domain controller.
- *Edb.log*: The Edb.log file is a transaction log. When changes occur to Active Directory objects, the changes are initially saved to the transaction log before they are written to the Active Directory database.
- *Edbxxxxx.log*: This is auxiliary transaction logs that can be used in cases where the primary Edb.log file fills up prior to it being written to the Ntds.dit Active Directory database.
- *Edb.chk*: Edb.chk is a checkpoint file that is used by the transaction logging process.
- *Res log files*: These are reserve log files whose space is used if insufficient space exists to create the Edbxxxxx.log file.
- *Temp.edb*: Temp.edb contains information on the transactions that are being processed.
- *Schema.ini*: The Schema.ini file is used to initialize the Ntds.dit Active Directory database when a domain controller is promoted.

What is LDAP?

The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that runs directly over the TCP/IP stack. The information model (both for data and namespaces) of LDAP is similar to that of the X.500 OSI directory service, but with fewer features and lower resource requirements than X.500. Unlike most other Internet protocols, LDAP has an associated API that simplifies writing Internet directory service applications. The LDAP API is applicable to directory management and browser applications that do not have directory service support as their primary function. LDAP cannot create directories or specify how a directory service operates.

Can you connect Active Directory to other 3rd-party Directory Services? Name a few options.

Where is the AD database held? What other folders are related to AD?

What is the SYSVOL folder?

The sysVOL folder stores the server's copy of the domain's public files. The contents such as group policy, users etc of the sysvol folder are replicated to all domain controllers in the domain.

The sysvol folder must be located on an NTFS volume

The article describes how to use the Burflags registry entry to rebuild each domain controller's copy of the system volume (SYSVOL) tree on all domain controllers in a common Active Directory directory service domain.

The term SYSVOL refers to a set of files and folders that reside on the local hard disk of each domain controller in a domain and that are replicated by the File Replication service (FRS). Network clients access the contents of the SYSVOL tree by using the following shared folders:

- NETLOGON
- SYSVOL

We recommend the procedure that is described in this article as a last resort to restore a domain's SYSVOL tree and its contents. Use this procedure only if you cannot make the FRS functional on individual domain controllers in the domain. Use this procedure only if the bulk restart can be performed more quickly than troubleshooting and resolving replication inconsistencies, and time to resolution is a critical factor.

Important Domain controllers will not service authentication request during the procedure. Only when the SYSVOL and NETLOGON folders are shared again will the domain controller authenticate requests. This procedure should not be performed during peak hours.

Note See the "How to temporarily stabilize the domain SYSVOL tree" section of this article for information about how to temporarily stabilize the domain SYSVOL tree until you can complete all the steps in the "How to rebuild the domain system volume replica set across enterprise environments" section.

We strongly recommend that you monitor FRS performance and health by using monitoring tools. By using monitoring tools, you may prevent the need for replica set authoritative and non-authoritative restores, and you may provide insight into the root cause of FRS failures. The following monitoring tool is available for download:

Name the AD NCs and replication issues for each NC

What are application partitions? When do I use them

Application Directory Partition is a partition space in Active Directory which an application can use to store that application specific data. This partition is then replicated only to some specific domain controllers.

The application directory partition can contain any type of data except security principles (users, computers, groups).

How do you create a new application partition

Create an application directory partition by using the DnsCmd command

Use the **DnsCmd** command to create an application directory partition. To do this, use the following syntax:

DnsCmd ***ServerName*** **/CreateDirectoryPartition** ***FQDN of partition***

To create an application directory partition that is named CustomDNSPartition on a domain controller that is named DC-1, follow these steps:

1. Click **Start**, click **Run**, type cmd, and then click **OK**.
2. Type the following command, and then press ENTER:

dnscmd DC-1 /createdirectorypartition CustomDNSPartition.contoso.com

When the application directory partition has been successfully created, the following information appears:
DNS Server DC-1 created directory partition: CustomDNSPartition.contoso.com Command completed successfully.

How do you view replication properties for AD partitions and DCs?

Configure Intersite Replication Availability

Updated: June 18, 2007

To control the blocks of time during which intersite replication can occur over a site link, you can use the Active Directory Sites and Services snap-in to configure the availability settings in the site link schedule.

Membership in the **Enterprise Admins** group in the forest or the **Domain Admins** group in the forest root domain, or equivalent, is the minimum required to complete this procedure. Review details about using the appropriate accounts and group memberships at http://go.microsoft.com/fwlink/?LinkId=83477.

To configure intersite replication availability

1. Open Active Directory Sites and Services. To open Active Directory Sites and Services, click **Start**, click **Administrative Tools**, and then click **Active Directory Sites and Services**.
2. In the console tree, click the intersite transport folder that contains the site link for which you are configuring intersite replication availability.

   Where?

   - Active Directory Sites and Services/Sites/Inter-Site Transports/IP or SMTP
3. In the details pane, right-click the site link whose schedule you want to configure, and then click **Properties**.
4. Click **Change Schedule**.

   Note

   When you are logged on with an account that does not have sufficient credentials to change the schedule, the available option is **View Schedule**.

5. Select the block of time during which you want replication to be either available or not available, and then click **Replication Not Available** or **Replication Available**, respectively.

Backup and Restore Tasks for Active Directory

Active Directory is backed up as part of system state, a collection of system components that depend on each other. You must back up and restore system state components together.

Components that comprise the system state on a domain controller include:

- **System Start-up Files (boot files).** These are the files required for Windows 2000 Server to start.
- **System registry.**
- **Class registration database of Component Services.** The Component Object Model (COM) is a binary standard for writing component software in a distributed systems environment.
- **SYSVOL.** The system volume provides a default Active Directory location for files that must be shared for common access throughout a domain. The SYSVOL folder on a domain controller contains:
    - NETLOGON shared folders. These usually host user logon scripts and Group Policy objects (GPOs) for non-Windows 2000based network clients.
    - User logon scripts for Windows 2000 Professionalbased clients and clients that are running Windows 95, Windows 98, or Windows NT 4.0.
    - Windows 2000 GPOs.
    - File system junctions.
    - File Replication service (FRS) staging directories and files that are required to be available and synchronized between domain controllers.
- **Active Directory.** Active Directory includes:
    - Ntds.dit: The Active Directory database.
    - Edb.chk: The checkpoint file.
    - Edb*.log: The transaction logs, each 10 megabytes (MB) in size.
    - Res1.log and Res2.log: Reserved transaction logs.

Note: If you use Active Directory-integrated DNS, then the zone data is backed up as part of the Active Directory database. If you do not use Active Directory-integrated DNS, you must explicitly back up the zone files. However, if you back up the system disk along with the system state, zone data is backed up as part of the system disk.If you installed Windows Clustering or Certificate Services on your domain controller, they are also backed up as part of system state. Details of these components are not discussed in this guide.

Top of page

General Guidelines for Backup

The backup tool in Windows 2000 Server supports multiple types of backup: *normal*, *copy*, *incremental*, *differential*, and *daily*. However, because Active Directory is backed up as part of system state, the only type of backup available for Active Directory is *normal*. A normal backup creates a backup of the entire system state while the domain controller is online. In addition, the backup tool marks each file as a backed up file, which clears the archive attribute of the file.

Considerations for ensuring a good backup

To ensure a successful restore from backup, you must know what defines a *good backup*.

Which domain controllers to back up

At a minimum, back up two domain controllers in each domain, one of which should be an operations master role holder (excluding the relative ID (RID) master, which should not be restored). Note that backup data from a domain controller can only be used to restore that domain controller. You cannot use a backup of one domain controller to restore another.

Contents

A good backup includes at least the system state and the contents of the system disk. Backing up the system disk ensures that all the required system files and folders are present so you can successfully restore the data.

Note: Best performance practice states that the Active Directory's logs and database files should be on separate disks. If you have configured your domain controllers in this manner you will have Active Directory components spread out on multiple drives, such as D:\Winnt\NTDS for your logs and E:\Winnt\NTDS for your database. You do not need to specify these log and database locations in order for them to be backed up; the backup utility will automatically locate and include them when you back up system state.

Age

A backup that is older than the tombstone lifetime set in Active Directory is not a good backup. At a minimum, perform at least two backups within the tombstone lifetime. The default tombstone lifetime is 60 days. Active Directory incorporates the tombstone lifetime into the backup and restore process as a means of protecting itself from inconsistent data.

Deleting an object from Active Directory is a two-step process. When an object is deleted in Active Directory, the object gets converted into a tombstone, which is then replicated to the other domain controllers in the environment to inform them of the deletion. Active Directory purges the tombstone when the tombstone lifetime is reached.

If you restore a domain controller to a state prior to the deletion of an object, and the tombstone for that object is not replicated to the restored domain controller before the tombstone expires, the object remains present only on the restored domain controller, resulting in inconsistent data. Thus, you must restore the domain controller prior to expiration of the tombstone, and allow inbound replication from a domain controller containing the tombstone to complete prior to expiration of the tombstone.

Active Directory protects itself from restoring data older than the tombstone lifetime by disallowing the restore. As a result, the useful life of a backup is equivalent to the tombstone lifetime setting for the enterprise.

Top of page

General Guidelines for Restore

You can start the restore process by using either the Windows 2000 Server backup utility or another supported utility. You can perform either a non-authoritative restore or an authoritative restore.

Top of page

How to Select the Appropriate Restore Method

You select the appropriate restore method by considering:

- Circumstances and characteristics of the failure. The two major categories of failure, from an Active Directory perspective, are Active Directory data corruption and hardware failure. Active Directory data corruption occurs when the directory contains

corrupt data that has been replicated to all domain controllers or when a large portion of the Active Directory hierarchy has been changed accidentally (such as deletion of an OU) and this change has replicated to other domain controllers.

- Roles and functions of the failed server.

Non-authoritative restore of Active Directory

A non-authoritative restore returns the domain controller to its state at the time of backup, then allows normal replication to overwrite that state with any changes that have occurred after the backup was taken. After you restore the system state, the domain controller queries its replication partners. The replication partners replicate any changes to the restored domain controller, ensuring that the domain controller has an accurate and updated copy of the Active Directory database.

Non-authoritative restore is the default method for restoring Active Directory, and you will use it in most situations that result from Active Directory data loss or corruption. To perform a non-authoritative restore, you must be able to start the domain controller in Directory Services Restore Mode.

Non-authoritative restore of SYSVOL

When you non-authoritatively restore the SYSVOL, the local copy of SYSVOL on the restored domain controller is compared with that of its replication partners. After the domain controller restarts, it contacts its replication partners, compares SYSVOL information, and replicate the any necessary changes, bringing it up-to-date with the other domain controllers within the domain.

Perform a non-authoritative restore of SYSVOL if at least one other functioning domain controller exists in the domain. This is the default method for restoring SYSVOL and occurs automatically if you perform a non-authoritative restore of the Active Directory.

If no other functioning domain controller exists in the domain, then perform a primary restore of the SYSVOL. A primary restore builds a new File Replication service (FRS) database by loading the data present under SYSVOL on the local domain controller. This method is the same as a non-authoritative restore, except that the SYSVOL is marked primary.

Authoritative restore of Active Directory

An authoritative restore is an extension of the non-authoritative restore process. You must perform the steps of a non-authoritative restore before you can perform an authoritative restore. The main difference is that an authoritative restore has the ability to increment the version number of the attributes of all objects in an entire directory, all objects in a subtree, or an individual object (provided that it is a leaf object) to make it authoritative in the directory. Restore the smallest unit necessary, for example, do not restore the entire directory in order to restore a single subtree.

As with a non-authoritative restore, after a domain controller is back online, it will contact its replication partners to determine any changes since the time of the last backup. However, because the version number of the object attributes that you want to be authoritative will be higher than the existing version numbers of the attribute held on replication partners, the object on the restored domain controller will appear to be more recent and therefore will be replicated out to the rest of the domain controllers within the environment.

Unlike a non-authoritative restore, an authoritative restore requires the use of a separate tool, Ntdsutil.exe. No backup utilities— including the Windows 2000 Server system tools— can perform an authoritative restore.

An authoritative restore will not overwrite new objects that have been created after the backup was taken. You can authoritatively restore only objects from the configuration and domain-naming contexts. Authoritative restores of schema-naming contexts are not supported.

Perform an authoritative restore when human error is involved, such as when an administrator accidentally deletes a number of objects and that change replicates to the other domain controllers and you cannot easily recreate the objects. To perform an authoritative restore, you must start the domain controller in Directory Services Restore Mode.

Authoritative restore of SYSVOL

By authoritatively restoring the SYSVOL, you are specifying that the copy of SYSVOL that is restored from backup is authoritative for the domain. After the necessary configurations have been made, Active Directory marks the local SYSVOL as authoritative and it is replicated to the other domain controllers within the domain.

The authoritative restore of SYSVOL does not occur automatically after an authoritative restore of Active Directory. Additional steps are required.

As with Active Directory authoritative restore, you typically perform an authoritative restore of SYSVOL when human error is involved and the error has replicated to other domain controllers. For example, you might perform an authoritative restore of SYSVOL if an administrator has accidentally deleted an object that resides in SYSVOL, such as a Group Policy object.

Recover a domain controller through reinstallation

To recover a domain controller through reinstallation, you do not restore the system state from backup media; instead, you reinstall Windows, install Active Directory, and allow replication partners to bring the recovered domain controller up to date.

Recovering a domain controller through reinstallation can quickly return the computer to service if the following conditions exist:

- A domain controller has failed and you cannot restart in Directory Services Restore mode. If failure was caused by a hardware failure, you have resolved the hardware problem (for example, by replacing the disk).
- There are other domain controllers in the domain, to serve as replication partners.
- The computer is functioning only as a domain controller (it does not run other server services such as Exchange), and it does not contain other data that needs to be recovered from a backup.

Restore a domain controller through reinstallation and restore from backup

This method involves first reinstalling Windows 2000, to enable you to start in Directory Services Restore Mode. During the Windows 2000 Server setup process, you will obtain more information about the nature of the failure and you can then determine whether you can reinstall Windows 2000 Server into the same partition as it was previously installed or

whether you will need to re-partition the drive. After you successfully reinstall Windows 2000, you can start in Directory Services Restore Mode and perform a normal non-authoritative restore from backup media.

Restore a domain controller through reinstallation and restore the system state from backup if the following conditions exist:

- A domain controller has failed and you cannot restart in Directory Services Restore mode. If failure was caused by a hardware failure, you have resolved the hardware problem (for example, by replacing the disk).
- You have the following information about the failed domain controller:
    - Disk configuration. You need a record of the volumes and sizes of the disks and partitions. You use this information to recreate the disk configuration in the case of a complete disk failure. You must recreate all disk configurations prior to restoring system state. Failure to recreate all disk configurations can cause the restore process to fail and can prevent you from starting the domain controller following the restore.
    - Computer name. You need the computer name to restore a domain controller of the same name and avoid changing client configuration settings.
    - Domain membership. You must know the domain name because even if the computer name does not change, you might need to re-establish a new computer account.
    - Local Administrator password. You must know the local computer's Administrator password that was used when the backup was created. Without it, you will not be able to log on to the computer to establish a domain account for the computer after you restore it. If you are not part of the domain, you will not be able to log on by using a domain account, even if you are a domain administrator. The local Administrator password is also required to restore the system state on a domain controller.
- The domain controller is running other server services such as Exchange, or contains other data you must restore from a backup.
- You have a good backup, made within the tombstone lifetime.

Considerations for restoring operations masters

To restore an operations master role holder, you must perform one of the following procedures:

- Restore the failed operations master from backup.
- Seize the role to another domain controller within the environment. Seize the operations master role only if you do not intend to restore the original role holder from backup. For more information about seizing operations master roles, see "Managing Operations Masters" in this guide.

Restoring the RID Master can result in Active Directory data corruption, so it is not recommended.

Restoring the Schema Master can result in orphaned objects, so it is not recommended.

Considerations for recovering global catalog servers

To recover the global catalog server you can either:

- Restore the failed global catalog server from backup.

- Assign a new global catalog to compensate for the loss of the original.

Restoring from backup is the only way that a domain controller that was functioning as a global catalog at the time of backup can automatically be restored to the role of global catalog. Restoring a domain controller by reinstallation does not automatically reinstate the global catalog role. In a multi-domain environment, be aware that restoring a global catalog server from backup requires more time than restoring a domain controller that does not host the global catalog.

As there are no real disadvantages in configuring multiple global catalogs, you might want to create a new global catalog in your environment if you anticipate an extended downtime for the failed global catalog server. Creating a new global catalog server is particularly relevant if users associated with the original global catalog server can no longer access a global catalog server, or if the requirement for the global catalog service is significant in your environment, such as when you are running Exchange 2000.

For more information about creating a new global catalog server, see "Managing Global Catalogs Servers" in this guide.

Note: Configuring multiple global catalogs servers in a forest increases the availability of the system, but also increases replication traffic and database size. If you do restore the failed domain controller and maintain its role as a global catalog server, you might want to remove any additional global catalogs servers that you configured during its absence.

Considerations for restoring onto different hardware

It is possible to restore a domain controller onto different hardware. However, you should consider the following issues:

- **Different hardware abstraction layers (HALs).** By default, the Hal.dll is not backed up as part of system state, however the Kernel32.dll is. Therefore, if you try to restore a backup onto a computer that requires a different HAL (for example, to support a multiprocessor environment) compatibility issues exist between the new HAL and the original Kernel32.dll. To overcome this incompatibility, manually copy the **Hal.dll** from the original computer and install it on the new computer. The limitation is that the new computer can use only a single processor.
- **Incompatible Boot.ini File.** If you backup and restore the boot.ini file, you might have some incompatibility with your new hardware configuration, resulting in a failure to start. Before you restore it, ensure that the boot.ini file is correct for your new hardware environment.
- **Different Network or Video Cards.** If your new hardware has a different video adapter or multiple network adapters, then uninstall them before you restore data. When you restart the computer; the normal Plug and Play functionality makes the necessary changes.
- **Disk Space and Partition Configuration**. Partitions on the new computer must match those on the original computer. Specifically, all the drive mappings must be the same and the partition size must be at least equal to that on the original computer.

Considerations for authoritative restores

Performing an authoritative restore can affect group membership and passwords for trusts and computer accounts.

Impact on group membership

By performing an authoritative restore, you risk possible loss of group membership information.

Because group membership is a multi-valued attribute, and because of how Active Directory handles links, back links and deletions, an authoritative restore can produce varying results to group membership. These variations are based on which objects replicate first after an authoritative restore: the User object or the Group object.

If the un-deletion of the user replicates first, then the group membership information of both the group (the members it contains) and the user (the groups to which the user belongs) will be represented correctly.

If the un-deletion of the group replicates first, the replication partners will drop the addition of the (locally) deleted user from the group membership. The only exception to this is the user's primary group, which is always represented correctly both from the user and group reference.

You cannot control which object replicates first after you perform an authoritative restore. If your environment is affected by this situation, the only option is to modify the group membership attribute of the affected groups on the domain controller where you performed the authoritative restore.

This issue stems not from the integrity of the restored data, but from the way in which the data is replicated. By looking at this domain controller, administrators can view the way the directory should look and take steps to replicate the accurate directory information to the other domain controllers within the domain.

The best way to do this is to add a fictitious user and then delete that same fictitious user to and from each group that was involved in the authoritative restore.

A group is involved in the restore if it was either authoritatively restored itself or if it had members restored who did not have that group defined as their primary group.

By doing this, you force the correct group membership information to be replicated out from the source domain controller (the domain controller on which you performed the original authoritative restore) and update the group membership information on its replication partners. These updated objects reflect the correct memberships and also correct the information represented in the **Member of** tab of the restored user objects' properties.

You must ensure that no additions are made to group membership (for the affected groups and users) on any of the other domain controllers within the environment.

If you do not adhere to this process, the accurate version of the directory (held on the domain controller where the restore was performed) can become corrupted by the incorrect membership information. If the accurate version of the directory becomes corrupted, you must either update group membership manually or perform another authoritative restore of the objects by using the **verinc** option, and perform the process again.

Impact on trusts and computer accounts

In Windows 2000, trust relationships and computer account passwords are negotiated at a specified interval (by default 30 days for trust relationships and computer passwords).

When you perform an authoritative restore, you might restore previously used passwords for the objects in the Active Directory that maintain trust relationships and computer accounts.

In the case of trust relationships, this can impact communication with other domain controllers from other domains, causing permissions errors when users try to access resources in other domain. To rectify this, you must remove and recreate NTLM trust relationships to Windows 2000 or Windows NT 4.0 domains.

In the case of a computer account password, this can impact communications between the member workstation or server and a domain controller of its domain. This effect might cause users on Windows NT or Windows 2000 computers to have authentication difficulty due to an invalid computer account.

Top of page

Backup and Restore Tasks and Procedures

Table 1.8 shows the tasks and procedures for backup and restore.

Table 1.8 Backup and Restore Tasks and Procedures

Tasks	Procedures	Tools	Frequency
Back up Active Directory and associated components.	• Back up system state on a domain controller. • Back up system state and system disk on a domain controller.	• NTBackup.exe	At least twice within the tombstone lifetime
Perform a non-authoritative restore.	• Restart the domain controller in Directory Services Restore Mode (locally or remotely). • Restore from backup media. • Verify Active Directory restore.	• NTBackup.exe • Ntdsutil.exe • Event Viewer • Repadmin.exe	As needed
Perform an authoritative restore of a subtree or leaf object.	Restart in Directory Services Restore Mode. Restore from backup media for authoritative restore. Restore system state to an alternate	NTBackup.exe Ntdsutil.exe Event Viewer Repadmin.exe	As needed

	location. Perform authoritative restore of the subtree or leaf object. Restart in normal mode. Restore applicable portion of SYSVOL from alternate location. Verify Active Directory restore.		
Perform an authoritative restore of the entire directory.	Restart in Directory Services Restore Mode. Restore from backup media for authoritative restore. Restore system state to an alternate location. Restore the database. Restart in normal mode. Copy SYSVOL from alternate location. Verify Active Directory restore.	NTBackup.exe Ntdsutil.exe Event Viewer Repadmin.exe	As needed
Recover a domain controller through reinstallation.	Clean up metadata. Install Windows 2000 Server. Install Active Directory.	Ntdsutil.exe Active Directory Sites and Services Active Directory Users and Computers Dcpromo.exe	As needed
Restore a domain controller through reinstallation and subsequent restore	Install Windows 2000 Server on the same drive letter and partition as before	• NTBackup.exe	As needed

from backup.	the failure, partitioning the drive if necessary. Restore from backup media (non-authoritative restore). Verify Active Directory restore.	

Top of page

What is the Global Catalog?

What Is the Global Catalog?

Updated: December 5, 2005

In this section

The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain Active Directory forest. The global catalog is stored on domain controllers that have been designated as global catalog servers and is distributed through multimaster replication. Searches that are directed to the global catalog are faster because they do not involve referrals to different domain controllers.

In addition to configuration and schema directory partition replicas, every domain controller in a Windows 2000 Server or Windows Server 2003 forest stores a full, writable replica of a single domain directory partition. Therefore, a domain controller can locate only the objects in its domain. Locating an object in a different domain would require the user or application to provide the domain of the requested object.

The global catalog provides the ability to locate objects from any domain without having to know the domain name. A global catalog server is a domain controller that, in addition to its full, writable domain directory partition replica, also stores a partial, read-only replica of all other domain directory partitions in the forest. The additional domain directory partitions are partial because only a limited set of attributes is included for each object. By including only the attributes that are most used for searching, every object in every domain in even the largest forest can be represented in the database of a single global catalog server.

Note

- A global catalog server can also store a full, writable replica of an application directory partition, but objects in application directory partitions are not replicated to the global catalog as partial, read-only directory partitions.

The global catalog is built and updated automatically by the Active Directory replication system. The attributes that are replicated to the global catalog are identified in the schema as the partial attribute set (PAS) and are defined by Microsoft. However, to optimize searching, you can edit the schema by adding or removing attributes that are stored in the global catalog.

In Windows 2000 Server environments, any change to the PAS results in full synchronization (update of all attributes) of the global catalog. Windows Server 2003 reduces the impact of updating the global catalog by replicating only the attributes that change.

In a single-domain forest, a global catalog server stores a full, writable replica of the domain and does not store any partial replica. A global catalog server in a single-domain forest functions in the same manner as a non-global-catalog server except for the processing of forestwide searches.

Common Global Catalog Scenarios

The following events require a global catalog server:

- Forestwide searches. The global catalog provides a resource for searching an Active Directory forest. Forestwide searches are identified by the LDAP port that they use. If the search query uses port 3268, the query is sent to a global catalog server.
- User logon. In a forest that has more than one domain, two conditions require the global catalog during user authentication:
    - In a Windows 2000 native mode domain or a Windows Server 2003 domain at either the Windows 2000 native or Windows Server 2003 domain functional level, domain controllers must request universal group membership enumeration from a global catalog server.
    - When a user principal name (UPN) is used at logon and the forest has more than one domain, a global catalog server is required to resolve the name.
- Universal Group Membership Caching: In a forest that has more than one domain, in sites that have domain users but no global catalog server, Universal Group Membership Caching can be used to enable caching of logon credentials so that the global catalog does not have to be contacted for subsequent user logons. This feature eliminates the need to retrieve universal group memberships across a WAN link from a global catalog server in a different site.

    Note

    Universal groups are available only in a Windows 2000 Server native mode domain or a Windows Server 2003 domain at either the Windows 2000 native or Windows Server 2003 domain functional level.

- Exchange Address Book lookups. Servers running Microsoft Exchange 2000 Server and Exchange Server 2003 rely on access to the global catalog for address information. Users use global catalog servers to access the global address list (GAL).

How do you view all the GCs in the forest?

Provide me any testcase for this, YOU should check GCs are in SYNC. no need to browse it, use ADSI TO get complete naming convention extracted from GC

Why not make all DCs in a large forest as GCs?

Trying to look at the Schema, how can I do that?

regsvr32 schmmgmt.**dll**

What is LDP?

LDP is a Lightweight Directory Access Protocol (LDAP) client utility that is included with Microsoft Windows 2000. This article describes the basics of how to query and browse an LDAP-compliant directory by using the LDP utility.

What is REPLMON

This article describes how to use the Active Directory Replication Monitor (ReplMon.exe) tool to determine the servers that hold the operations master roles in a forest as well as the domain controllers and global catalog servers for the forest. The five operations master roles that are defined in Active Directory are:

- Schema master
- Domain naming master
- Relative identifier (RID) master
- Primary domain controller (PDC) emulator
- Infrastructure master

There is a very quick method to determine which servers in the forest hold these roles by using Active Directory Replication Monitor.

What is ADSIEDIT

When a new user is created in Active Directory, the **Full name** field is always generated in FirstName LastName format. In turn, this field sets the **Display Name** field on creation, therefore, you end up with a FirstName LastName formatted global address list.

You can make this change by using the Adsiedit utility. Adsiedit not only changes the default way the **Display Name** field is built, but also the **Full Name** (that is, the "cn") field, therefore, users appear in the chosen format when you look in the Users and Computers snap-in.

ADSIEdit Instructions

Warning If you use the ADSI Edit snap-in, the LDP utility, or any other LDAP version 3 client, and you incorrectly modify the attributes of Active Directory objects, you can cause serious problems. These problems may require you to reinstall Microsoft Windows 2000 Server, Microsoft Windows Server 2003, Microsoft Exchange 2000 Server, Microsoft Exchange Server 2003, or both Windows and Exchange. Microsoft cannot guarantee that

problems that occur if you incorrectly modify Active Directory object attributes can be solved. Modify these attributes at your own risk.

1. Insert your Windows 2000 Server CD.
2. Navigate to the **\support\tools** directory.
3. Double-click on the **Support.cab** file.
4. Locate the files adsiedit.msc and adsiedit.dll. Extract them to your **%systemroot%\system32** directory.
5. Run **regsvr32 adsiedit.dll**.
6. Start Microsoft Management Console (MMC), and then add the ADSI Edit snap-in.
7. Right-click the top node, and then click **Connect to**.
8. Change the **Naming Context** to "Configuration Container," and then click **OK** to bind and authenticate.
9. Expand the **Configuration Container** node, and then expand the **Configuration** node.
10. Expand the **cn=DisplaySpecifiers** node, and then double-click **CN=409**.**NOTE**: 409 is the Locale ID for U.S. English. If you are in a multi-lingual environment, you may need to make changes to the other codes. Most of the Asian codes are already set.

    The International Telecommunication Union (ITU) and International Organization for Standardization (ISO) define the code pages. For more information, visit the following ITU and ISO Web pages:

    http://www.itu.int

    http://www.microsoft.com/globaldev/reference/lcid-all.mspx

    For more information about supporting localized Exchange Clients, click the following article number to view the article in the Microsoft Knowledge Base:

    150977 Supporting localized Exchange clients

11. In the right-hand pane, open the properties for "CN=user-Display".
12. Scroll to the **createDialog** optional property.
13. Set the attribute to %<*sn*>.%<*givenName*>. Make sure that you click **Set**.

    Note The only tokens that can be formatted in the dislayName are %<sn>, %<givenName>, and %<initials>.
14. Click **OK** to close the dialog box.

15. In Active Directory Users and Computers, create a new User; the Full Name (and thus, the Display Name) are built in accordance with your rule.

What is NETDOM

You can use the Netdom.exe tool to reset the secure channel between a workstation, server, or domain controller. This article describes the syntax for variations that you can use with Netdom.exe.

What is Repadmin

Repadmin.exe is a Microsoft Windows 2000 Resource Kit tool that is available in the Support Tools folder on the Windows 2000 CD-ROM. It is a command-line interface to Active Directory replication. This tool provides a powerful interface into the inner workings of Active Directory replication, and is useful for troubleshooting Active Directory replication problems. This article describes the basic use of the Repadmin.exe tool.

What are sites? What are they used for?

- An Activity Directory site is a region of a network that has high bandwidth connectivity.

- A site is a collection of one or more subnets connected by high speed links.
- A site may span multiple domains.
- A domain may span multiple sites.

Sites are used for the following functions:

- To optimize replication for speed and bandwidth consumption between domain controllers

- To locate the closest domain controller for client logon, services, and directory searches
- To direct DFS client to the server in the site
- To optimize the replication of Sysvol

What's the difference between a site link's schedule and interval?

Description of a Site

A site is a collection of one or more subnets that are defined by the administrator. When you define subnets, they should be "well-connected" with high-bandwidth local area network (LAN) connections.

Sites can contain multiple domains, and a domain can span more than one site. If a domain spans more than one site, it must replicate by using the Internet Protocol (IP) inter-site transport. You can use the Simple Mail Transfer Protocol (SMTP) inter-site transport only for global catalog replication and replication of non-domain naming contexts, such as the configuration and schema.

You define and administer a site in the "Active Directory Sites and Services Manager" snap-in. When you install a domain controller as the first domain controller in a forest, a new site is created by default. You can also create other sites.

Description of a Site Link / A site link is an object that typically represents two sites that are connected physically by a wide area network (WAN) link. Although the site link may contain more than two sites, this article discusses the simplest case--a link that represents two sites.

The site link allows the administrator to assign the cost and transport for replication. This procedure defines parameters for replication. The cost is an arbitrary value that is selected by the administrator to reflect the speed and reliability of the physical connection between the sites. When you lower the cost value on the link, the priority is increased. Site links have a replication interval and a schedule that are independent of the cost. The cost is used by the KCC to prefer one site link path over another.

If a site link has more than two sites, all of the sites in the site link are considered connected in a NxN fully connected star topology.

The KCC uses site links to decide which sites to link with connections. Without site links, the KCC has no information about the sites that are reachable on the network and does not know the relative costs of the WAN links between the sites. You should add at least enough site links so that every site is transitively linked to every other site. When you do this, a directory object that is added or modified on a particular domain controller in a particular site eventually makes its way to all of the domain controllers in all of the sites.

What is the KCC?

The Knowledge Consistency Checker (KCC) is an Active Directory component that is responsible for the generation of the replication topology between domain controllers. This article describes the role of one server per site, known as the Inter-Site Topology Generator, which is responsible for managing the inbound replication connection objects for all bridgehead servers in the site in which it is located.

What is the ISTG? Who has that role by default?

Connection objects for bridgehead servers for inter-site replication are created differently. The KCC on one domain controller (regardless of the domain) in each site is responsible for reviewing the inter-site topology and creating inbound replication connection objects as necessary for bridgehead servers in the site in which it resides. This domain controller

is known as the Inter-Site Topology Generator (ISTG). The domain controller holding this role may not necessarily also be a bridgehead server.

When the ISTG determines that a connection object needs to be modified on a given bridgehead server in the site, the ISTG makes the change to its local Active Directory copy. As part of the normal intra-site replication process, these changes propagate to the bridgehead servers in the site. When the KCC on the bridgehead server reviews the topology after receiving these changes, it translates the connection objects into replication links that Active Directory uses to replicate data from remote bridgehead servers.

What are the requirements for installing AD on a new server?

- An NTFS partition with enough free space
- An Administrator's username and password
- The correct operating system version
- A NIC
- Properly configured TCP/IP (IP address, subnet mask and - optional - default gateway)
- A network connection (to a hub or to another computer via a crossover cable)
- An operational DNS server (which can be installed on the DC itself)
- A Domain name that you want to use
- The Windows 2000 or Windows Server 2003 CD media (or at least the i386 folder)
- Brains (recommended, not required...)

How can you forcibly remove AD from a server, and what do you do later? • Can I get user passwords from the AD database?

Dcpromo /forceremoval, an administrator can forcibly remove Active Directory and roll back the system without having to contact or replicate any locally held changes to another DC in the forest.

Reboot the server then After you use the dcpromo /forceremoval command, all the remaining metadata for the demoted DC is not deleted on the surviving domain controllers, and therefore you must manually remove it by using the NTDSUTIL command. In the event that the NTDS Settings object is not removed correctly you can use the Ntdsutil.exe utility to manually remove the NTDS Settings object.

You will need the following tool: Ntdsutil.exe, Active Directory Sites and Services, Active Directory Users and Computers

Name some OU design considerations.

What is tombstone lifetime attribute?

The tombstone lifetime must be substantially longer than the expected replication latency between the domain controllers. The interval between cycles of deleting tombstones must be at least as long as the maximum replication propagation delay across the forest. Because the expiration of a tombstone lifetime is based on the time when an object was deleted logically, rather than on the time when a particular server received that tombstone through replication, an object's tombstone is collected as garbage on all servers at approximately the same time. If the tombstone has not yet replicated to a particular domain controller, that DC never records

the deletion. This is the reason why you cannot restore a domain controller from a backup that is older than the tombstone lifetime.

How would you find all users that have not logged on since last month?

What are the DS* commands?

The DS (Directory Service) group of commands are split into two families. In one branch are DSadd, DSmod, DSrm and DSMove and in the other branch are DSQuery and DSGet.

DS Syntax

These DS tools have their own command structure which you can split into five parts:

1 2 3 4 5
Tool object "DN" (as in LDAP distinguished name) -switch value For example:
DSadd user "cn=billy, ou=managers, dc=cp, dc=com" -pwd cX49pQba

This will add a user called Billy to the Managers OU and set the password to cx49Qba

Here are some of the common DS switches which work with DSadd and DSmod
-pwd (password) -upn (userPrincipalName) -fn (FirstName) -samid (Sam account name).

What's the difference between LDIFDE and CSVDE? Usage considerations?

This step-by-step article describes how to use the Csvde.exe utility to create contacts and user accounts in Active Directory. You may have to use this method in some scenarios, for example, when administrators want to export custom recipients from Microsoft Exchange Server 5.5 and import them into Active Directory as Microsoft Windows contacts.

Csvde.exe is a Microsoft Windows 2000 command-line utility that is located in the SystemRoot\System32 folder after you install Windows 2000. Csvde.exe is similar to Ldifde.exe, but it extracts information in a comma-separated value (CSV) format. You can use Csvde to import and export Active Directory data that uses the comma-separated value format. Use a spreadsheet program such as Microsoft Excel to open this .csv file and view the header and value information. See Microsoft Excel Help for information about functions such as **Concatenate** that can simplify the process of building a .csv file.

Note Although Csvde is similar to Ldifde, Csvde has a significant limitation: it can only import and export Active Directory data by using a comma-separated format (.csv). Microsoft recommends that you use the Ldifde utility for Modify or Delete operations. Additionally, the distinguished name (also known as DN) of the item that you are trying to import must be in

the first column of the .csv file or the import will not work.

The source .csv file can come from an Exchange Server directory export. However, because of the difference in attribute mappings between the Exchange Server directory and Active Directory, you must make some modifications to the .csv file. For example, a directory export from Exchange Server has a column that is named "obj-class" that you must rename to "objectClass." You must also rename "Display Name" to "displayName."

For more information about attribute mappings, click the following article number to view the article in the Microsoft Knowledge Base:
281563 Exchange Server 5.5 to Exchange 2000 attribute mappings for the Migration Wizard
Use the following syntax to run the tool from a command prompt:
csvde -i -f c:\filename.csv

What are the FSMO roles? Who has them by default? What happens when each one fails?

Flexible Single Master Operation (FSMO) role. Currently in Windows 2000 there are five FSMO roles:

- Schema master
- Domain naming master
- RID master
- PDC emulator
- Infrastructure daemon

Back to the top

Schema Master FSMO Role

The schema master FSMO role holder is the DC responsible for performing updates to the directory schema (that is, the schema naming context or LDAP://cn=schema,cn=configuration,dc=<domain>). This DC is the only one that can process updates to the directory schema. Once the Schema update is complete, it is replicated from the schema master to all other DCs in the directory. There is only one schema master per directory.

Back to the top

Domain Naming Master FSMO Role

The domain naming master FSMO role holder is the DC responsible for making changes to the forest-wide domain name space of the directory (that is, the Partitions\Configuration

naming context or LDAP://CN=Partitions, CN=Configuration, DC=<domain>). This DC is the only one that can add or remove a domain from the directory. It can also add or remove cross references to domains in external directories.

Back to the top

RID Master FSMO Role

The RID master FSMO role holder is the single DC responsible for processing RID Pool requests from all DCs within a given domain. It is also responsible for removing an object from its domain and putting it in another domain during an object move.

When a DC creates a security principal object such as a user or group, it attaches a unique Security ID (SID) to the object. This SID consists of a domain SID (the same for all SIDs created in a domain), and a relative ID (RID) that is unique for each security principal SID created in a domain.

Each Windows 2000 DC in a domain is allocated a pool of RIDs that it is allowed to assign to the security principals it creates. When a DC's allocated RID pool falls below a threshold, that DC issues a request for additional RIDs to the domain's RID master. The domain RID master responds to the request by retrieving RIDs from the domain's unallocated RID pool and assigns them to the pool of the requesting DC. There is one RID master per domain in a directory.

Back to the top

PDC Emulator FSMO Role

The PDC emulator is necessary to synchronize time in an enterprise. Windows 2000 includes the W32Time (Windows Time) time service that is required by the Kerberos authentication protocol. All Windows 2000-based computers within an enterprise use a common time. The purpose of the time service is to ensure that the Windows Time service uses a hierarchical relationship that controls authority and does not permit loops to ensure appropriate common time usage.

The PDC emulator of a domain is authoritative for the domain. The PDC emulator at the root of the forest becomes authoritative for the enterprise, and should be configured to gather the time from an external source. All PDC FSMO role holders follow the hierarchy of domains in the selection of their in-bound time partner.

In a Windows 2000 domain, the PDC emulator role holder retains the following functions:

- Password changes performed by other DCs in the domain are replicated preferentially to the PDC emulator.

- Authentication failures that occur at a given DC in a domain because of an incorrect password are forwarded to the PDC emulator before a bad password failure message is reported to the user.
- Account lockout is processed on the PDC emulator.
- The PDC emulator performs all of the functionality that a Microsoft Windows NT 4.0 Server-based PDC or earlier PDC performs for Windows NT 4.0-based or earlier clients.

This part of the PDC emulator role becomes unnecessary when all workstations, member servers, and domain controllers that are running Windows NT 4.0 or earlier are all upgraded to Windows 2000. The PDC emulator still performs the other functions as described in a Windows 2000 environment.

The following information describes the changes that occur during the upgrade process:

- Windows 2000 clients (workstations and member servers) and down-level clients that have installed the distributed services client package do not perform directory writes (such as password changes) preferentially at the DC that has advertised itself as the PDC; they use any DC for the domain.
- Once backup domain controllers (BDCs) in down-level domains are upgraded to Windows 2000, the PDC emulator receives no down-level replica requests.
- Windows 2000 clients (workstations and member servers) and down-level clients that have installed the distributed services client package use the Active Directory to locate network resources. They do not require the Windows NT Browser service.

Back to the top

Infrastructure FSMO Role

When an object in one domain is referenced by another object in another domain, it represents the reference by the GUID, the SID (for references to security principals), and the DN of the object being referenced. The infrastructure FSMO role holder is the DC responsible for updating an object's SID and distinguished name in a cross-domain object reference.

NOTE: The Infrastructure Master (IM) role should be held by a domain controller that is not a Global Catalog server(GC). If the Infrastructure Master runs on a Global Catalog server it will stop updating object information because it does not contain any references to objects that it does not hold. This is because a Global Catalog server holds a partial replica of every object in the forest. As a result, cross-domain object references in that domain will not be updated and a warning to that effect will be logged on that DC's event log.

If all the domain controllers in a domain also host the global catalog, all the domain

controllers have the current data, and it is not important which domain controller holds the infrastructure master role.

What FSMO placement considerations do you know of?

FSMO availability and placement

Dcpromo.exe performs the initial placement of roles on domain controllers. This placement is often correct for directories with few domain controllers. In a directory with many domain controllers the default placement is unlikely to be the best match to your network.

On a per-domain basis, select local primary and standby FSMO domain controllers in case a failure occurs on the primary FSMO owner. Additionally, you may want to select off-site standby owners in the event of a site-specific disaster scenario. Consider the following in your selection criteria:

- If a domain has only one domain controller, that domain controller holds all the per-domain roles.
- If a domain has more than one domain controller, use Active Directory Sites and Services Manager to select direct replication partners with persistent, "well-connected" links.
- The standby server may be in the same site as the primary FSMO server for faster replication convergence consistency over a large group of computers, or in a remote site in the event of a site-specific disaster at the primary location.
- Where the standby domain controller is in a remote site, ensure that the connection is configured for continuous replication over a persistent link.

Back to the top

General recommendations for FSMO placement

- Place the RID and PDC emulator roles on the same domain controller. It is also easier to keep track of FSMO roles if you cluster them on fewer machines.

  If the load on the primary FSMO load justifies a move, place the RID and primary domain controller emulator roles on separate domain controllers in the same domain and active directory site that are direct replication partners of each other.
- As a general rule, the infrastructure master should be located on a nonglobal catalog server that has a direct connection object to some global catalog in the forest, preferably in the same Active Directory site. Because the global catalog server holds a partial replica of every object in the forest, the infrastructure master, if placed on a global catalog server, will never update anything, because it does not contain any

references to objects that it does not hold. Two exceptions to the "do not place the infrastructure master on a global catalog server" rule are:

- Single domain forest:

  In a forest that contains a single Active Directory domain, there are no phantoms, and so the infrastructure master has no work to do. The infrastructure master may be placed on any domain controller in the domain, regardless of whether that domain controller hosts the global catalog or not.
- Multidomain forest where every domain controller in a domain holds the global catalog:

  If every domain controller in a domain that is part of a multidomain forest also hosts the global catalog, there are no phantoms or work for the infrastructure master to do. The infrastructure master may be put on any domain controller in that domain.

- At the forest level, the schema master and domain naming master roles should be placed on the same domain controller as they are rarely used and should be tightly controlled. Additionally, the domain naming master FSMO should also be a global catalog server. Certain operations that use the domain naming master, such as creating grand-child domains, will fail if this is not the case.

  In a forest at the Forest Functional Level Windows Server 2003, you do not have to place the domain naming master on a global catalog.

Most importantly, confirm that all FSMO roles are available using one of the management consoles (such as Dsa.msc or Ntdsutil.exe).

I want to look at the RID allocation table for a DC. What do I do?

What's the difference between transferring a FSMO role and seizing one? Which one should you NOT seize? Why?

How do you configure a "stand-by operation master" for any of the roles?

A standby operations master is a domain controller that you identify as the computer that assumes the operations master role if the original computer fails. A single domain controller can act as the standby operations master for all of the operations master roles in a domain, or you can designate a separate standby for each operations master role.

No utilities or special steps are required to designate a domain controller as a standby operations master. However, the current operations master and the standby should be well connected. This means that the network connection between them must support at least a 10-megabit transmission rate and be available at all times. In addition, configure the current role

holder and the standby as direct replication partners by manually creating a Connection object between them.

Configuring a replication partner can save some time if you must reassign any operations master roles to the standby operations master. Before transferring a role from the current role holder to the standby operations master, ensure that replication between the two computers is functioning properly. Because they are replication partners, the new operations master is as updated as the original operations master, thus reducing the time required for the transfer operation.

During role transfer, the two domain controllers exchange any unreplicated information to ensure that no transactions are lost. If the two domain controllers are not direct replication partners, a substantial amount of information might need to be replicated before the domain controllers completely synchronize with each other. The role transfer requires extra time to replicate the outstanding transactions. If the two domain controllers are direct replication partners, fewer outstanding transactions exist and the role transfer operation completes sooner.

Designating a domain controller as a standby also minimizes the risk of role seizure. By making the operations master and the standby direct replication partners, you reduce the chance of data loss in the event of a role seizure, thereby reducing the chances of introducing corruption into the directory.

When you designate a domain controller as the standby, follow all recommendations that are discussed in Guidelines for Role Placement in Introduction to Administering Operations Master Roles. To designate a standby for the forest-level roles, choose a global catalog server so it can interact more efficiently with the domain naming master. To designate a standby for the domain-level roles, ensure that the domain controller is not a global catalog server so that the infrastructure master continues to function properly if you must transfer the roles.

Task Requirements

The following tools are required to perform the procedures for this task:

- Active Directory Sites and Services
- Repadmin.exe

How do you backup AD?

Backing up Active Directory is essential to maintain the proper health of the Active Directory database. You can backup Active Directory by using the NTBACKUP tool that comes built-in with Windows Server 2003, or use any 3rd-party tool that supports this feature. Backing up the Active Directory is done on one or more of your Active Directory domain Controllers (or DCs), and is performed by backing up the System State on those servers. The System State contains the local Registry, COM+ Class Registration Database, the System Boot Files, certificates from Certificate Server (if it's installed), Cluster database (if it's installed), NTDS.DIT, and the SYSVOL folder

To ensure your ability to actually use this backup, you must be aware of the tombstone lifetime. By default, the tombstone is 60 days (for Windows 2000/2003 DCs), or 180 days (for Active Directory based upon Windows Server 2003 SP1 DCs).

Longer tombstone lifetime decreases the chance that a deleted object remains in the local directory of a disconnected DC beyond the time when the object is permanently deleted from online DCs. The tombstone lifetime is not changed automatically when you upgrade to Windows Server 2003 with SP1, but you can change the tombstone lifetime manually after the upgrade. New forests that are installed with Windows Server 2003 with SP1 have a default tombstone lifetime of 180 days. Read my "Changing the Tombstone Lifetime Attribute in Active Directory" article for more info on that.

Any backup older than 60/180 days is not a good backup and cannot be used to restore any DC. You do not need to backup all your DCs' System States, usually backing up the first DC in the Forest + the first DCs in each domain is enough for most scenarios.

Purpose of Performing Regular Backups

You need a current, verified, and reliable backup to:

- Restore Active Directory data that becomes lost. By using an authoritative restore process, you can restore individual objects or sets of objects (containers or directory partitions) from their deleted state. Read my "Recovering Deleted Items in Active Directory" article for more info on that.
- Recover a DC that cannot start up or operate normally because of software failure or hardware failure.
- Install Active Directory from backup media (using the **dcpromo /adv** command). Read my "Install DC from Media in Windows Server 2003" article for more info on that.
- Perform a forest recovery if forest-wide failure occurs.

All these are reasons to have good working and reliable backups.

Note: One of the Active Directory features that was introduced in Windows Server 2003 with Service Pack 1 was the Directory Service Backup Reminders. With this reminder, a new event message, event ID 2089, provides the backup status of each directory partition that a domain controller stores. This includes application directory partitions and Active Directory Application Mode (ADAM) partitions. If halfway through the tombstone lifetime a partition has not been backed up, this event is logged in the Directory Service event log and continues daily until the partition is backed up.

Note: You can only back up the System State data on a local computer. You cannot back up the System State data on a remote computer.

How do you restore AD?

In the Windows Server 2003 family, you can restore the Active Directory database if it becomes corrupted or is destroyed because of hardware or software failures. You must restore the Active Directory database when objects in Active Directory are changed or deleted.

Note: There is an option to restore Active Directory objects that have been deleted and are now in a phase called "tombstone". These items are hidden from the GUI and await their cleanup by a process called "garbage collection". Read more about it on my "Recovering Deleted Items in Active Directory" article.

You can use one of the three methods to restore Active Directory from backup media: Primary Restore, Normal Restore (i.e. Non Authoritative), and Authoritative Restore.

Primary Restore: This method rebuilds the first domain controller in a domain when there is no other way to rebuild the domain. Perform a primary restore only when all the domain controllers in the domain are lost, and you want to rebuild the domain from the backup. Members of the Administrators group can perform the primary restore on local computer. On a domain controller, only members of the Domain Admins group can perform this restore.

Normal Restore: This method reinstates the Active Directory data to the state before the backup, and then updates the data through the normal replication process. Perform a normal restore for a single domain controller to a previously known good state.

Authoritative Restore: You perform this method in tandem with a normal restore. An authoritative restore marks specific data as current and prevents the replication from overwriting that data. The authoritative data is then replicated through the domain. Perform an authoritative restore for individual object in a domain that has multiple domain controllers. When you perform an authoritative restore, you lose all changes to the restore object that occurred after the backup. You need to use the NTDSUTIL command line utility to perform an authoritative restore. You need to use it in order to mark Active Directory objects as authoritative, so that they receive a higher version recently changed data on other domain controllers does not overwrite System State data during replication.

For example, if you inadvertently delete or modify objects in Active Directory, and those objects were thereafter replicated to other DCs, you will need to authoritatively restore those objects so they are replicated or distributed to the other servers. If you do not authoritatively restore the objects, they will never get replicated or distributed to your other servers because they will appear to be older than the objects currently on your other DCs. Using the NTDSUTIL utility to mark objects for authoritative restore ensures that the data you want to restore gets replicated or distributed throughout your organization.

On the other hand, if your system disk has failed or the Active Directory database is corrupted, then you can simply restore the data normally without using NTDSUTIL. After rebooting the DC, it will receive newer updates from other DCs.

How do you change the DS Restore admin password?

1. If you know the password for the offline administrator account, start the recovery domain controller in Dsrepair mode. If you do not know the password for the offline administrator account, reset the password while the recovery domain controller is still in normal Active Directory mode.

   You can use the **setpwd** command-line tool to reset the password on domain controllers that are running Microsoft Windows 2000 Service Pack 2 (SP2) and later

while they are in online Active Directory mode.

Note Microsoft no longer supports Windows 2000 SP2. Install the most recent Windows 2000 service pack to obtain this functionality.

For more information about changing the Recovery Console administrator password, click the following article number to view the article in the Microsoft Knowledge Base:

239803 How to change the Recovery Console administrator password on a domain controller

Administrators of Windows Server 2003 domain controllers can use the **set dsrm password** command in the **Ntdsutil** command-line tool to reset the password for the offline administrator account.

For more information about how to reset the Directory Services Restore Mode administrator account, click the following article number to view the article in the Microsoft Knowledge Base:

322672 How to reset the Directory Services Restore Mode administrator account password in Windows Server 2003

About AD Backup

Active Directory is a hierarchical database that holds information about the network's resources such as computers, servers, users, groups and more. The main purpose of Active Directory is to provide central authentication and authorization services. Normal administrative tasks when working with Active Directory include creating, managing, moving, editing and sometimes – deleting – various objects such as user accounts, computer accounts, groups, contacts and other objects. The Active Directory database is stored on Domain Controllers (or DCs), in a file called NTDS.DIT (that's not everything, but it'll do for a short intro…)

While deleting an object in Active Directory is usually something an administrator would think twice before doing, sometimes mistakes do happen, and then the administrator ends up with one (or more) deleted items that he or she cannot restore anymore.

How does Active Directory treat deleted items?

When an object is deleted from Active Directory, it is not immediately erased, but is marked for future deletion. You see, Active Directory uses a replication model that is characterized as "multi-master loose consistency with convergence". Changes can be made on any DC in the forest, and the changes are then incrementally replicated throughout the forest. Therefore, object deletions in this environment cannot simply remove an object, because doing so would remove the unit of replication itself.

The marker used to designate that an AD object scheduled to be destroyed is called "tombstone". A tombstone is an object whose **IsDeleted** property has be set to **True**, and it indicates that the object has been deleted but not removed from the directory, much like a deleted file is removed from the file allocation table but the data is not actually removed from the drive. The directory service moves tombstoned objects to the Deleted Objects container, where they remain until the garbage collection process removes the objects. The garbage collection process by default runs every 12 hours on a DC. The length of time tombstoned objects remain in the directory service before being deleted is either 60 days for Windows 2000/2003 Active Directory, or 180 days for Windows Server 2003 SP1 Active Directory (by default). The tombstone lifetime must be significantly longer than the garbage collection frequency to ensure that deletion of objects is replicated to other DCs.

Considering all the above, a delete operation is essentially a special modify operation that:

1. Sets the *IsDeleted* value to True.
2. Sets the internal *WhenDeleted* column to the *IsDeleted* metadata's *TimeChanged* time stamp.
3. Sets the Windows NT security descriptor to a special value.
4. Changes the relative distinguished name (RDN) to a value that is otherwise impossible, (that is, one that cannot be set by an LDAP program).
5. Strips all attributes not needed at this point by Active Directory. Key attributes such as the following are hard-coded to survive deletion:
    - Object-GUID
    - Object-SID
    - Object-Dist-Name
    - USN

Note: You can make changes to the Active Directory that allow the survival of more attributes in case of an object deletion. This was covered in our article entitled - Protect Objects in Windows Server 2003 Active Directory from Accidental Deletion.

You must understand the difference between restoring an object that has long been deleted from the database, and no longer is present in it, not even as a tombstoned object, and restoring a tombstoned object. Restoring tombstoned objects from the Active Directory database is often known as "reanimation", and this is what this article is about.

Because tombstoning an object strips it from many attributes, you must know that if you do elect to reanimate a deleted user or group, you will still have to recover the group memberships and any other linked attributes of which you might be in need. Also, without going too deep into this issue, know that you cannot reanimate objects that were deleted from the Configuration NC (or Partition). I will try to cover these issues in a future article.

Note: One of the Active Directory features that were introduced in Windows Server 2003 with Service Pack 1 was the Directory Service Backup Reminders. With this reminder, a new event message, event ID 2089, provides the backup status of each directory partition that a domain controller stores, including application directory partitions and Active Directory Application Mode (ADAM) partitions. If halfway through the tombstone lifetime a partition has not been backed up, this event is logged in the Directory Service event log and continues daily until the partition is backed up.

Methods for restoring deleted items in Active Directory

There are several methods of reanimating tombstoned objects from the Active Directory. Some are simple and easy to perform, some are more cumbersome. Some are freeware, some

are more sophisticated and cost (a lot of) money. On this page I've listed some of the freely available tools. For those that cost money – hire a consultant (or me…).

Whatever you do, make sure you have a good and working backup of the domain controller's System State. The System State contains the local Registry, COM+ Class Registration Database, the System Boot Files, certificates from Certificate Server (if it's installed), Cluster database (if it's installed), NTDS.DIT, and the SYSVOL folder.

You can easily backup the DC's System State by using NTBACKUP or any range of 3rd-party tools that have that capability built in them. You do not need to backup all your DCs' System States, usually backing up the first DC in the Forest + the first DCs in each domain is enough for most scenarios. You can read more about it on my "Backup Windows Server 2003 Active Directory" article (insert link).

Restoring the item from a previous backup

Restoring deleted items from a previous System State backup is not as simple as it sounds. In fact, this is not really reanimation, but actually a total restore of the deleted object. However, since restoring deleted items by usage of the NTBACKUP program and the System State backup involve shutting down the DC and booting it into "DS Restore Mode", the reanimation mechanism is the only way to recover deleted objects without taking a DC offline.

There are several issues and steps that you need to perform, all are covered in my "Restore Windows Server 2003 Active Directory" article

Restoring the objects with LDP.EXE

As written in the beginning of this article, deleted objects in Active Directory are not really deleted, they are just "tombstoned" for a period of time that can either be 60/180 days, depending on your DCs' operating system, or any other value, if it was ever changed by the system administrator.

Restoring objects with ADRestore.net

Guy Teverovsky, a fellow MVP from Israel, has written a cool tool that allows you to easily restore deleted AD objects. The tool is provided as freeware and has no kind of support, but from what I've seen, it works great. Some of the tools features include:

- Browsing the tombstones
- Domain Controller targeting
- Can be used with alternative credentials (convenient if you do not logon to your desktop as Domain Admin, which you should never do anyway)
- User/Computer/OU/Container reanimation
- Preview of tombstone attributes

Enumerating tombstones

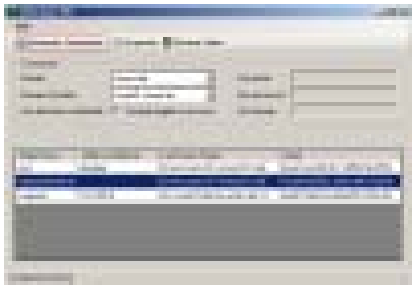

Previewing the tombstone attributes

Restoring a deleted user account

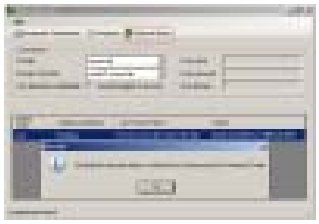

Download ADRestore.net

For more information on Guy's tool, please see Guy's blog entry announcing ADRestore.net

Restoring objects with Microsoft ADRestore (previously Sysinternals)

Formerly Sysinternals and now Microsoft, Mark Russinovich has created a command-line freeware application called ADRestore. The tool enumerates all of the currently tombstoned objects in a domain and allows you to restore them selectively, and provides a convenient command-line interface for using the Active Directory reanimation functionality. If you run it from the command line you will be prompted to choose which object you want to restore, and since there could be quite a few tombstoned objects, this process might take some time as you answer NO to each and every prompt.

To add a little selectivity to the restore operation, you can run ADRestore with a parameter to narrow down the search. For example

would search for all objects with "daniel" as part of its name.

The -r switch forces the program to prompt the user for each restoration. Otherwise, all the objects found matching said criteria will be automatically restored. The default (no criteria supplied) is that all tombstoned objects will be enumerated and restored.

Note that deleted items may no longer be members of specific organizational units or OUs. Restoring these objects from deleted status will not automatically restore them to their respective OUs; this will need to be done manually.

Download ADRestore

How to restore deleted user accounts and their group memberships in Active Directory - 840001

Quest Object Restore for Active Directory

Quest Software, a leading provider of application, database and Windows management solutions, offers at no charge a graphical utility that helps Microsoft Active Directory administrators recover deleted objects using the Tombstone Reanimation feature of Windows Server 2003. This Microsoft recovery interface allows administrators to restore accidentally deleted objects online, without rebooting a domain controller. Quest Object Restore for Active Directory enhances this ability by providing a graphical interface, similar to the Windows Recycle Bin, for viewing and restoring Active Directory objects.

Restoring single, deleted objects in Active Directory can be a manual and time-consuming process requiring system downtime. Object Restore for Active Directory is a free, graphical utility that allows you to instantly recover deleted objects in a Windows Server 2003 environment without rebooting a Domain Controller. The freeware utility allows viewing Tombstoned objects in Active Directory and reanimating deleted items using Microsoft's new Tombstone Reanimation interfaces for Windows Server 2003. When you download the Freeware, a 6-month key is built in. You will be prompted to re-register on our site at the end of each 6-month period.

In order to download their product you will need to go through a very nagging and unfriendly registration screen. Proceed from here:

Quest: Object Restore for Active Directory

Note that Quest has a great variety of tools for Active Directory management and recovery, however since they are not freeware I will not give them a free advertising ride…

-

Why can't you restore a DC that was backed up 4 months ago?

Plz refer above AD backup and restore question and answers.

What are GPOs?

Group Policy objects

Policy settings are stored in Group Policy objects (GPOs). Settings for each GPO are edited using the Group Policy Object Editor. After installation of the Group Policy Management Console (GPMC), Group Policy Object Editor is usually opened from GPMC. For information about Group Policy Object Editor, see Group Policy object editor overview for GPMC. For information about GPMC, see Group Policy Management Console Overview.

There are two kinds of GPOs:

- **Active Directory-based GPOs**. These are stored in a domain and replicate to all the domain controllers for the domain. They are available only in an Active Directory environment. They apply to users and computers in a site, domain, or organizational unit to which the Group Policy object is linked. This is the primary mechanism through which Group Policy is used in an Active Directory environment.
- **Local GPOs**. There is just one local GPO stored on each computer. Local GPOs are the least influential GPOs in an Active Directory environment, and local GPOs have only a subset of the settings found in Active Directory-based GPOs. For information about local GPOs, see Local Group Policy objects overview for GPMC.

What is the order in which GPOs are applied?

Name a few benefits of using GPMC.

What are the GPC and the GPT? Where can I find them?

GPO version number in the GPC

I have used one of two tools to look at the raw GPO version number in the GPC; **Adsiedit** and **Ldp**. Both are GUI tools that allow you to look at Active Directory objects. I'm going to give an example using Ldp.exe simply because I can display all the relevant information in a single frame for the purposes of this discussion. Before going on, I would like to say a word of caution here. Making changes using low-level Active Directory editing tools could cause problems with the functionality of your domain. If you want to look at Active Directory objects for learning purposes, try this out on a test domain (that cannot harm VPs in any way).

I open the Ldp tool, connect, and bind it to my test domain. From the **view** menu, I select **tree**. From the root of the domain, I navigate through the system container to expand the policies container, where the GPC portion of my GPOs is stored. In the below picture, I have highlighted a sample GPO I called test1 which is identified by a unique GUID and has a version number of 262146. Using the procedure described earlier, I figured out this version number is equal to a computer version of 2 and a user version of 4.

The GPT is stored in the sysvol portion of the domain controllers' file system; for example, for my test domain called corp.fourthcoffee.com the GPT is stored at \\corp.fourthcoffee.com\sysvol\corp.fourthcoffee.com\Policies. Each GPO is stored under this

folder using the GUID. As you can see from the picture above, the GPC also defines an attribute, gPCFileSysPath, which contains the sysvol path to the GPO's GPT. In my example, the full path to my test1 GPO is \\corp.fourthcoffee.com\sysvol\corp.fourthcoffee.com\Policies\{06C1CDF4-7288-4A6D-B887-8727C2823857}. And I can then type the contents of the GPT.INI file stored directly under that folder to see the version number, as shown in the picture below. Notice this version number (262146) matches the number we saw for the GPC in the Active Directory. That is what we would expect.

GPC – Group Policy Container - The GPC is the store of the GPOs; The GPC is where the GPO stores all the AD-related configuration. Any GPO that is created is not effective until it is linked to an OU, Domain or a Site. The GPOs are replicated among the Domain Controllers of the Domain through replication of the Active Directory.

What are GPO links? What special things can I do to them?

To apply the settings of a GPO to the users and computers of a domain, site, or OU, you need to add a link to that GPO. You can add one or more GPO links to each domain, site, or OU by using GPMC. Keep in mind that creating and linking GPOs is a sensitive privilege that should be delegated only to administrators who are trusted and understand Group Policy

What can I do to prevent inheritance from above?

Policy inheritance

In general, Group Policy is passed down from parent to child containers within a domain, which you can view by using Active Directory Users and Computers. Group Policy is not inherited from parent to child domains, for example, from wingtiptoys.com to sales.wingtiptoys.com. Active Directory Domains and Trusts, which you can use to manage relationships of this type, is not related to Group Policy.

If you assign a specific Group Policy setting to a high-level parent container, that Group Policy setting applies to all containers beneath the parent container, including the user and computer objects in each container. However, if you explicitly specify a Group Policy setting for a child container, the child container's Group Policy setting overrides the parent container's setting.

If a parent organizational unit has policy settings that are not configured, the child organizational unit does not inherit them. Policy settings that are disabled are inherited as disabled. In addition, if a policy setting is configured (enabled or disabled) for a parent organizational unit and the same policy setting is not configured for a child organizational unit, the child inherits the parent's enabled or disabled policy setting.

If a policy setting that is applied to a parent organizational unit and a policy setting that is applied to a child organizational unit are compatible, the child organizational unit inherits the parent policy setting, and the child's setting is also applied.

If a policy setting that is configured for a parent organizational unit is incompatible with the same policy setting that is configured for a child organizational unit (because the setting is enabled in one case and disabled in the other), the child does not inherit the policy setting from the parent. The policy setting in the child is applied.

Blocking inheritance

You can block policy inheritance at the domain or organizational-unit level by opening the properties dialog box for the domain or organizational unit and selecting the **Block Policy inheritance** check box. For more information, see Block policy inheritance.

Enforcing inheritance

You can enforce policy inheritance by setting the **No Override** option on a Group Policy object link.

When you select the **No Override** check box, you force all child policy containers to inherit the parent's policy, even if that policy conflicts with the child's policy and even if **Block Inheritance** has been set for the child.

You can set **No Override** on a Group Policy object link by opening the

How can I override blocking of inheritance?

Name a few differences in Vista GPOs

- The Group Policy infrastructure
- Improved network awareness
- Added Group Policy capabilities

all take place under the hood. What administrators will find, however, is that Windows Vista™ Group Policy is much more powerful than it was in previous versions.
Prior to Windows Vista, Group Policy processing occurred within a process called winlogon. Winlogon had a lot of responsibility, which included getting people logged on to their desktops, as well as servicing the various Group Policy chores. Group Policy is now its own Windows® service. What's more, it's hardened, which means that it cannot be stopped nor can an administrator take ownership of the permissions upon Group Policy in order to then turn it off. These changes enhance the overall reliability of the Group Policy engine.
This is just for starters. Let's take a more in-depth look at some of the major changes that have been made to the new Group Policy

What are administrative templates?

Administrative Templates

Administrative templates, (or .adm files), enable administrators to control registry settings using Group Policy. These settings appear under the **Administrative Templates** folder for both user configuration and computer configuration in the console tree of the Group Policy Object Editor, and in HTML reports produced by GPMC.

It is important to understand that .adm files are not the actual settings that are deployed to client operating systems. The .adm file is simply a template file (implemented as text file with an .adm extension) that provides the friendly name for the setting and an explanation. This template file is used to populate the user interface. The settings that are deployed to clients are contained in the registry.pol file inside the GPO. On Windows XP and Windows Server 2003, each registry setting contains a "Supported on" tag that indicates which operating system versions support that policy setting. If a setting is specified and deployed to a client operating system that does not support that setting, the settings are ignored. These .adm files are stored in two locations by default: inside GPOs, and in the %windir%\inf folder on the local computer.

Windows includes a predefined set of Administrative template files that define the registry settings that can be configured in a Group Policy object (GPO). The .adm files can be added or removed from the Group Policy Object Editor by right-clicking **Administrative Templates** and clicking **Add/Remove Templates**. Adding or removing .adm files does not affect which policies are processed by the Group Policy engine. It only affects whether a specific Administrative Template policy setting is displayed in the Group Policy Object Editor. For example, if you removed all the .adm files from the GPO via the **Add/Remove Templates** dialog box, no Administrative Template policy settings would be displayed under the Administrative Templates node. This will not affect the policies already stored in the Registry.pol file.

Administrative Template	Description
System.adm	System settings
Inetres.adm	Internet Explorer settings
Wmplayer.adm	Windows Media Player settings. This tool is not available on Windows XP 64-Bit Edition and the 64-bit versions of the Windows Server 2003 family.
Conf.adm	NetMeeting settings. This tool is not available on Windows XP Professional, 64-Bit Edition and the 64-bit versions of the Windows Server family.
Wuau.adm	Windows Update settings.

What's the difference between software publishing and assigning?

Because you can publish software for users, assign software to users, or assign software to computers, you can establish a workable combination of those three options to meet your software management goals. The following is a comparison of these methods.

Publishing software for users

Typically, after you publish a software package to users in a site, domain, or OU, the users can use **Add or Remove Programs** to install the software. An exception is when you publish an application in a new GPO, and you must simultaneously link the GPO to the users in a site, domain, or OU. If you link a GPO and deploy the software at the same time, you must refresh the Group Policy before the application appears in **Add or Remove Programs**. Additionally, the application can be installed by opening an associated document if the application is deployed to do that (if **Auto-Install** is selected).

The user can remove the software, and then later choose to reinstall it, by using **Add or Remove Programs**.

Assigning software to users

There are three methods for assigning software: assign to users on-demand, assign to users, or assign to computers.

Important

- Check software license agreements before you assign applications to users because assigning software can result in an application being installed on multiple computers.

Issues might occur, regardless of whether you use the policy setting option **Remove the application if it falls out of the scope of management.**

Note : for part II Exchange Server Technical Interview Questions and answer will be update soon.

Regards

Ajit Khot IT Eng

Kentz Co Ltd.
Kingdom of Saudi Arabia
Email : ajitskhot@gmail.com
Email: Ajit.khot@kentz.com
www.kentz.com