1. **What is Active Directory?**
   Ans: Microsoft's directory database for Windows 2000/2003 networks. Stores information about resources on the network and provides centrally organizing, managing, and controlling access to the resources.

2. **What is LDAP?**
   Ans: Lightweight Directory Access Protocol. It is a database of active directory and is used to store the active directory objects in windows 2000. It is named as Active Directory in windows 2000/2003.

3. **What is DNS?**
   Ans: Domain Name System) - The Internet naming scheme which consists of a hierarchical sequence of names, from the most specific to the most general (left to right), separated by dots,

   And it is the system which translates the internet domain name into IP address and vice-versa.

   The Server, which translates such types of request, is DNS server.

4. **How do you check DNS is working or how do you check the service record of DNS is working?**
   Go to command prompt

   After you have setup your DNS Server, it's very important to check that the entries which are populated to the Internet are correct. You can use the following checklist using **nslookup.**

   Start nslookup for the desired DNS Server

   **Nslookup**
   **> server 193.247.121.196**
   Default Server: rabbit.akadia.ch
   Address: 193.247.121.196

5. **What is DHCP?**
   Dynamic Host Configuration Protocol (DHCP) is a communications protocol that lets network administrators manage and automate the assignment of Internet Protocol (IP) addresses in an organization's network. DHCP allows devices to connect to a network and be automatically assigned an IP address.

6. **What is WINS? And what is the difference between WINS and DNS.**
   Windows Internet Name Server, it translates the NetBIOS name to IP address.

   DNS translate the FQDN name to IP address. FQDN Name consists of zone name, domain name and host name and these are separated by dots. And it resolves this address into IP address.

   WINS does not translate from FQDN name to IP address.

7. **Why DNS is required for AD not WINS?**

   The naming structure of Active Directory objects is based on Internet Naming System. It consists of hierarchal naming structure separated by dots. And to resolve computer record/ service record / mail exchange records, a service is required which support such type of translation/ resolution. And DNS fits very much to this service. In fact Active Directory is of no use without DNS. That's why DNS is very much required for Active directory.

8. **What is Disaster Recovery?**

   Disaster recovery is the process to bring the server on line in short period of time and less effect to the business in disaster.

   Disaster recovery is consists of

   a. Backup
   b. Recovery console
   c. ASR ( Automated System Recovery) in Windows 2003 and ERD in Windows 2000

**9. What is the difference between Authoritative and Non-Authoritative Restore?**

a. Authoritative Restore:

The main purpose of Authoritative restore is to undo or roll back changes that have been made to active directory, or to reset data stored in a distributed directory such as sysvol.

b. Non-Authoritative Restore.

The data and distributed services on a domain controller are restored from a backup media and then updated through normal replication.

Example: If a restore backup contains a user named "Mark" and the user was deleted after last backup, the mark user object will also be deleted on the restored domain controller via the replication process

Reason for Non-Authoritative Restore

I. Restoring a single domain controller in an environment that includes multiple domain controller

II. Attempting to restore SYSVOL or File Replication Service data on domain controllers.

c. Primary Restore

New in windows 2003

Reason for Primary Restore

I. Restoring the only domain controller in an Active Directory Environment

II. Restoring the first of several domain controllers

III. Restoring the first domain controller in a replica set.

When All the domain controller or the only domain controller in a domain have failed, primary restore in needed. If a domain is lost, the first domain controller should be restored as primary restore, and any subsequent domain controller should be restored using a Normal or Non-Authoritative restore.

**10. What is File Replication Service**

The replication service maintains identical sets of files and directories on different servers and workstations. When files are updated on one server, the file replication service replaces the corresponding files on other servers and workstations with the updated files. The replication process simplifies the task of updating and coordinating files, and maintains the integrity of the replicated data.

**11. What are backed up in System State Backup?**

In System State Backup

I. The System Registry

II. The COM+ Class Registration Database

III. The boot files

a) Boot.ini

b) Ntdetect.com

c) Ntldr

d) Bootsect.dos

e) Ntbootdd.sys

IV. System files protected by windows File Protection Service

V. Certificate Service Database if installed.

VI. Active Directory and Sysvol folder on the DC.

VII. Cluster Service information on cluster server

VIII. Internet Information Server Metabase.

**12. What is Forest, Domain, Schema, Global Catalogue, Universal Group Caching?**

Forest:

Tree is collection of hierarchal structure of domains that share a common name space and are connected by transitive trust relationship. And the collection of such trees is called forest.

Domain:

Active Directory environment are logical groupings of resources that ultimately forms units of replication. And logical grouping of these units are Domains.

Schema:

Schema represents the definitions of all objects types that exists within Active Directory and their associated attributes. The schema is stored on all domain controllers throughout the forest. It controls all updates and modification to schema. Schema is make up of class and attributes.

Global Catalogue:

I. It Information about all Active Directory Objects from all domains in a single forest.

II. Stores information of universal groups and their associated membership.

III. Forwarding authentication request to the appropriate domain when user principal name is used to log on.

IV. Validate object references within a forest.

Universal Group Membership Caching:

UGMC helps to reduce the number of universal group membership queries that need to be forwarded across a WAN link when a user attempts to log on.

By default UGMC updates the Universal Group Membership information every 8 hours for a user.

Benefit of UGMC:

I. Faster user login times, because the global catalogue server does not need to be contacted for all logon requests.

II. Reducing the need to place global catalogue servers in each site.

III. Reducing the usage of WAN bandwidth usage associated with Global Catalogue replication.

Note: Where a high member of directory queries are expected global catalogue at each site represents the best possible solution.

**13. What is stub zone?**

Stub zone is new in Windows 2003 Server. It contains read-only resource record which it obtains from other name servers. But it contains only three types of resource record

I. A copy of SOA record for the servers

II. Copies of NS records for all name servers authoritative for the zone.

III. Copies of A records for all name servers authoritative for the zone

It does not contain CNAME records, MX records, SRV records, or A records for other hosts in the zone. **The most important benefit for stub zone is to reduce the network traffic over WAN link connection and time to resolve the resource records queries.**

**14. What are the steps and commands to restore?**

I. Start the computer in Directory Service Restore Mode.

II. Restore the backup from the media as in Non-Authoritative Mode.

III. After restore complete do not restart the computer.

IV. Go to command prompt

V. Enter into ntdsutil

E:\ntdsutil>**ntdsutil**
ntdsutil: **authoritative restore**
authoritative restore: **restore object OU=bosses,DC=ourdom,DC=com**

Opening DIT database... Done.

The current time is 06-17-05 12:34.12.
Most recent database update occurred at 06-16-05 00:41.25.
Increasing attribute version numbers by 100000.

Counting records that need updating...
Records found: 0000000012

After Completing the authoritative restore restart the computer.

**15. How may types of backup and advantage and disadvantage?**

**Types of Backup**
There are different kinds of backups, the following lists some of them:
**Full Backup**
Full backup is the starting point for all other backups, and contains all the data in the folders and files that are selected to be backed up. Because full backup stores all files and folders, frequent full backups result in faster and simpler restore operations. Remember that when you choose other backup types, restore jobs may take longer.

| | |
|---|---|
| Advantages | Restore is the fastest |
| Disadvantages | Backing up are the slowest |
| | The storage space requirements are the highest |

**Incremental Backup**
An incremental backup backs up only those files created or changed since the last normal or incremental backup. It marks files as having been backed up (in other words, the archive attribute is cleared). If you use a combination of normal and incremental backups, you will need to have the last normal backup set as well as all incremental backup sets in order to restore your data.

| | |
|---|---|
| Advantages | Backing up are the fastest |
| | The storage space requirements are the lowest |
| Disadvantages | Restore is the slowest |

**Differential Backup**
A differential backup copies file created or changed since the last normal or incremental backup. It does not mark files as having been backed up (in other words, the archive attribute is not cleared). If you are performing a combination of normal and differential backups, restoring files and folders requires that you have the last normal as well as the last differential backup.

| | |
|---|---|
| Advantages | Restore is faster than restoring from incremental backup |
| | Backing up is faster than a full backup |
| | The storage space requirements are lower than for full backup |
| Disadvantages | Restore is slower than restoring from full backup |
| | Backing up is slower than incremental backup |
| | The storage space requirements are higher than for incremental backup |

**16. What are the ports numbers of these?**

| TCP/UDP Port | Number | Secure | Number |
|---|---|---|---|
| DNS | 53 | N/A | 53 |
| DHCP | 67 & 68 | N/A | 67 & 68 |
| RDP | 3389 | N/A | 3389 |
| LDAP | 389 | 636 | 389 |
| SMTP | 25 | 465 | 25 |
| POP | 110 | 995 | 110 |
| IMAP | 143 | 993 | 143 |
| HTTP | 80 | 443 | 80 |
| TELNET | 23 | 992 | 23 |
| KERBEROS | 88 | N/A | 88 |
| SNMP | 161 | N/A | 161 |
| IRC | 194 | N/A | 194 |
| NNTP | 119 | 563 | 119 |
| MS EXCH ROUTING | 691 | | 691 |
| MS SQL | 1433 | | 1433 |
| Ssh | 22 | | 22 |
| ftp | 21 | | 21 |

**17. How to know the size of Active directory Database?**

The size of ntds.dit will often be different sizes across the domain controllers in a domain. Remember that Active Directory is a multi-master independent model where updates are occurring in each of the ADs with the changes being replicated over time to the other domain controllers. The changed data is replicated between domain controllers, not the database, so there is no guarantee that the files are going to be the same size across all domain controllers.

Start/Reboot
Press F8
Choose Directory Services Restore Mode and press ENTER.
Press ENTER again to start the boot process.
Logon using the password defined for the local Administrator account
Open the Command Prompt
At the command prompt,
Run the ntdsutil command.
When ntdsutil has started
Type files and press ENTER.
Type info and then press ENTER. This will display current information about the path and size of the Active Directory database and its log files

**18. What are the files of active directory**

| | |
|---|---|
| NTDS.DIT | Main Database File |
| Edb.chk | Checkpoint log file |
| Edb.log | Transaction Log Files (If more than one log files, its name becomes edbhhhhhh.log. Where hhhhhh is the hexadecimal numbers. |
| .pat | Patch files - Manages data while backups are done. |
| res1.log & res2.log | Reserve log files - Reserves hard drive space for transaction log files |

**19. What is the location of NTDS.DIT file?**

%system root%\ntds\ntds.dit

**20. What is difference between Primary and Secondary DNS servers?**

Primary DNS server contains the read-write copy of the zone data base
Secondary DNS server contains read only copy of the zone database.
Primary DNS server can be an Active Directory Integrated.
Secondary DNS server can not be.

Secondary DNS server updates the DNS records from the Primary DNS server.

**21. How the Secondary DNS servers get the updates from Primary DNS server? Tell me the process**

The actual data transfer process is started by the client on client server mechanism.
With every new record entry, edit or update in the primary server, the serial number increases and it makes two changes one to the record and other to the zone serial number

The first record updated is

Record.
Others are not any specified order
The end of update is signaled by SOA record

**22. What is conditional forwarding in DNS?**

Conditional forwarding is the process to forward the client request to the exact DNS server for a particular domain name resolution request. DNS server needs to configure for conditional forwarding.
For example a DNS server is configured for domain.com
And a client sends a request for name resolution for host.microsoft.com. And this DNS server does not host any record for Microsoft.com domain. In this case this DNS server may be configured to forward all the name resolution request for Microsoft.com to the DNS server which hosts such types of records.
http://www.windowsnetworking.com/articles_tutorials/DNS_Conditional_Forwarding_in_Windows_Server_2003.html
To configure conditional forwarding
Go to DNS server properties
Forwarders
Click New
Enter the domain name in DNS domain text box
OK
Add the IP address for that domain
OK

**23. What is TCP and IP?**

TCP:

1. It provides a reliable-connection oriented packet delivery service
2. Guarantees delivery of IP datagram.
3. Perform segmentation and reassembly of large block of data sent by programs.
4. Ensures proper sequencing and ordered delivery of segmented data.
5. Perform check on the integrity of transmitted data by using checksum calculation
6. Sends acknowledgement of the received data.

IP:

1. IP is a connectionless, unreliable datagram protocol.
2. Primarily responsible for addressing and routing packets between hosts
3. IP does not attempt to recover from these types of error..

**A)The Schema Partition**

The schema partition stores two types of information: class and attribute schema definitions. The schema classes define all the types of objects that can be created and stored in Active Directory. The schema attributes define all the properties that can be used to describe the objects that are stored in Active Directory. When you install the first Exchange 2007 server role in the forest, the Active Directory preparation process adds many classes and attributes to the Active Directory schema. The classes that are added to the schema are used to create Exchange-specific objects,

such as agents and connectors. The attributes that are added to the schema are used to configure both the Exchange-specific objects, and the mail-enabled users and groups. These attributes include properties, such as Office Outlook Web Access settings and unified messaging settings. Every domain controller and global catalog server in the forest contains a complete replica of the schema partition.

**B) The Configuration Partition**

The configuration partition stores information about the forest-wide configuration. This data includes how Active Directory sites are configured, how information is displayed, and network services. Each type of configuration information is stored in a container in the configuration partition. Exchange configuration information is stored in a subfolder under the configuration partition's Services container. The information that is stored in this container includes the following:

- Address lists
- Address and display templates
- Administrative groups
- Connections
- Messaging Records Management, mobile, and unified messaging mailbox policies
- Global settings
- Recipient policies
- System policies
- Transport settings

Every domain controller and global catalog server in the forest contains a complete replica of the configuration partition.

**C) The Domain Partition**

The domain partition stores information in default containers and in organizational units that have been created by the Active Directory administrator. These containers hold the domain-specific objects. This data includes Exchange system objects and information about the as computers, users, and groups in that domain. When Exchange 2007 is installed, Exchange updates the objects in this partition to support Exchange functionalities. This includes storing and accessing recipient information. Each domain controller contains a complete replica of the domain partition for the domain for which it is authoritative. Every global catalog server in the forest contains a subset of every domain partition in the forest.

6) What is Global catalog ---?

The global catalog contains a partial replica of every Windows 2000 domain in the directory. The GC lets users and applications find objects in an Active Directory domain tree given one or more attributes of the target object. It also contains the schema and configuration of directory partitions. This means the global catalog holds a replica of every object in the Active Directory, but with only a small number of their attributes. ...

**Q3. How DNS really works**

DNS uses a client/server model in which the DNS server maintains a static database of domain names mapped to IP addresses. The DNS client, known as the resolver, performs queries against the DNS servers. The bottom line? DNS resolves domain names to IP address using these steps

Step 1. A client (or "resolver") passes its request to its local name server. For example, the URL term www.idgbooks.com typed into Internet Explorer is passed to the DNS server identified in the client TCP/IP configuration. This DNS server is known as the local name server.

Step 2. If, as often happens, the local name server is unable to resolve the request, other name servers are queried so that the resolver may be satisfied.

Step 3. If all else fails, the request is passed to more and more, higher-level name servers until the query resolution process starts with far-right term (for instance, com) or at the top of the DNS tree with root name servers

**Q4. Which are the major records in DNS?**

**1. Host or Address Records (A):-** map the name of a machine to its numeric IP address. In clearer terms, this record states the hostname and IP address of a certain machine. Have three fields: Host Name, Domain, Host IP Address.

**E.g.:-** eric.foobarbaz.com. IN A 36.36.1.6

It is possible to map more than one IP address to a given hostname. This often happens for people who run a firewall and have two Ethernet cards in one machine. All you must do is add a second A record, with every column the same save for the IP address.

**2. Aliases or Canonical Name Records (*CNAME*)**

"CNAME" records simply allow a machine to be known by more than one hostname. There must always be an A record for the machine before aliases can be added. The host name of a machine that is stated in an A record is called the canonical, or official name of the machine. Other records should point to the canonical name. Here is an example of a CNAME:

www.foobarbaz.com. IN CNAME eric.foobarbaz.com.

You can see the similarities to the previous record. Records always read from left to right, with the subject to be queried about on the left and the answer to the query on the right. A machine can have an unlimited number of CNAME aliases. A new record must be entered for each alias.

You can add A or CNAME records for the service name pointing to the machines you want to **load balance.**

**3. Mail Exchange Records (MX)**

MX" records are far more important than they sound. They allow all mail for a domain to be routed to one host. This is exceedingly useful – it abates the load on your internal hosts since they do not have to route incoming mail, and it allows your mail to be sent to any address in your domain even if that particular address does not have a computer associated with it. For example, we have a mail server running on the fictitious machine eric.foobarbaz.com. For convenience sake, however, we want our email address to be "user@foobarbaz.com" rather than "user@eric.foobarbaz.com". This is accomplished by the record shown below:

foobarbaz.com. IN MX 10 eric.foobarbaz.com.

The column on the far left signifies the address that you want to use as an Internet email address. The next two entries have been explained thoroughly in previous records. The next column, the number "10", is different from the normal DNS record format. It is a signifier of priority. Often larger systems will have backup mail servers, perhaps more than one. Obviously, you will only want the backups receiving mail if something goes wrong with the primary mail server. You can indicate this with your MX records. A lower number in an MX record means a higher priority, and mail will be sent to the server with the lowest number (the lowest possible being 0). If something happens so that this server becomes unreachable, the computer delivering the mail will attempt every other server listed in the DNS tables, in order of priority.

Obviously, you can have as many MX records as you would like. It is also a good idea to include an MX record even if you are having mail sent directly to a machine with an A record. Some sendmail programs only look for MX records.

It is also possible to include wildcards in MX records. If you have a domain where your users each have their own machine running mail clients on them, mail could be sent directly to each machine. Rather than clutter your DNS entry, you can add an MX record like this one:

*.foobarbaz.com. IN MX 10 eric.foobarbaz.com.

This would make any mail set to any individual workstation in the foobarbaz.com domain go through the server eric.foobarbaz.com.

One should use caution with wildcards; specific records will be given precedence over ones containing wildcards.

**4. Pointer Records (PTR)**

Although there are different ways to set up PTR records, we will be explaining only the most frequently used method, called "in-addr.arpa".

In-addr.arpa PTR records are the exact inverse of A records. They allow your machine to be recognized by its IP address. Resolving a machine in this fashion is called a "reverse lookup". It is becoming more and more common that a machine will do a reverse lookup on your machine before allowing you to access a service (such as a World Wide Web page). Reverse lookups are a good security measure, verifying that your machine is exactly who it claims to be. In-addr.arpa records look as such:

6.1.36.36.in-addr.arpa. IN PTR eric.foobarbaz.com.

As you can see from the example for the A record in the beginning of this document, the record simply has the IP address in reverse for the host name in the last column.

A note for those who run their own name servers: although Allegiance Internet is capable of pulling zones from your name server, we cannot pull the inverse zones (these in-addr.arpa records) unless you have been assigned a full class C network. If you would like us to put PTR records in our name servers for you, you will have to fill out the online web form on the support.allegianceinternet.com page.

**5. Name Server Records (NS)**

NS records are imperative to functioning DNS entries. They are very simple; they merely state the authoritative name servers for the given domain. There must be at least two NS records in every DNS entry. NS records look like this:

Foobarbaz.com. IN NS draven.foobarbaz.com.

There also must be an A record in your DNS for each machine you enter as A NAME server in your domain.

If Allegiance Internet is doing primary and secondary names service, we will set up these records for you automatically, with "nse.algx.net" and "nsf.algx.net" as your two authoritative name servers.

**6. Start Of Authority Records (SOA)**

The "SOA" record is the most crucial record in a DNS entry. It conveys more information than all the other records combined. This record is called the start of authority because it denotes the DNS entry as the official source of information for its domain. Here is an example of a SOA record, then each part of it will be explained:

foobarbaz.com. IN SOA draven.foobarbaz.com. hostmaster.foobarbaz.com. (

1996111901 ; Serial

10800 ; Refresh

3600 ; Retry

3600000 ; Expire

86400 ) ; Minimum

The first column contains the domain for which this record begins authority for. The next two entries should look familiar. The “draven.foobarbaz.com” entry is the primary name server for the domain. The last entry on this row is actually an email address, if you substituted a “@” for the first “.”. There should always be a viable contact address in the SOA record.

The next entries are a little more unusual then what we have become used to. The serial number is a record of how often this DNS entry has been updated. Every time a change is made to the entry, the serial number must be incremented. Other name servers that pull information for a zone from the primary only pull the zone if the serial number on the primary name server’s entry is higher than the serial number on it’s entry. In this way the name servers for a domain are able to update themselves. A recommended way of using your serial number is the YYYYMMDDNN format shown above; where the NN is the number of times that day the DNS has been changed.

Also, a note for Allegiance Internet customers who run their own name servers: even if the serial number is incremented, you should still fill out the web form and use the comment box when you make changes asking us to pull the new zones.

All the rest of the numbers in the record are measurements of time, in seconds. The “refresh” number stands for how often secondary name servers should check the primary for a change in the serial number. “Retry” is how long a secondary server should wait before trying to reconnect to primary server if the connection was refused. “Expire” is how long the secondary server should use its current entry if it is unable to perform a refresh, and “minimum” is how long other name servers should cache, or save, this entry.

**There can only be one SOA record per domain. Like NS records, Allegiance Internet sets up this record for you if you are not running your own name server.**

**Quick Summary of the major records in DNS**

| ***Record Type*** | ***Definition*** |
|---|---|
| Host (A) | Maps host name to IP address in a DNS zone. Has three fields: Domain, Host Name, Host IP Address. |
| Aliases (CNAME) | Canonical name resource record that creates an alias for a host name. CNAME records are typically used to hide implementation details from clients. Fields include: Domain, Alias Name, For Host DNS Name. |
| Name servers (NS) | Identifies the DNS name servers in the DNS domain. NS records appear in all DNS zones and reverse zones. Fields include: Domain, Name Server DNS Name. |
| Pointer (PTR) | Maps IP address to host name in a DNS reverse zone. Fields include: IP Address, Host DNS Name. |
| Mail Exchange (MX) | Specifies a mail exchange server for a DNS domain name. Note that the term “exchange” does not refer to Microsoft Exchange, a BackOffice e-mail application. However, to connect Microsoft Exchange to the Internet via the Internet Mail Server (IMS), the MX record must be correctly configured by your ISP.<br><br>A mail exchange server is a host that wills either process or forward mail for the DNS domain name. Processing the mail means either delivering it to the addressee or passing it to a different type of mail transport. Forwarding the mail means sending it to its final destination server, sending it using Simple Mail Transfer Protocol to another mail server that is closer to the final destination, or queuing it for a specified amount of time.<br><br>Fields include: Domain, Host Name (Optional), Mail Exchange Server DNS Name, Preference Number. |

**Q5.What is a DNS zone**

A zone is simply a contiguous section of the DNS namespace. Records for a zone are stored and managed together. Often, subdomains are split into several zones to make manageability easier. For example, *support.microsoft.com* and *msdn.microsoft.com* are separate zones, where *support* and *msdn* are subdomains within the Microsoft.com domain.

**Q6. Name the two Zones in DNS?**

DNS servers can contain *primary* and *secondary* zones. A primary zone is a copy of a zone where updates can be made, while a secondary zone is a copy of a primary zone. For fault tolerance purposes and load balancing, a domain may have several DNS servers that respond to requests for the same information.

The entries within a zone give the DNS server the information it needs to satisfy requests from other computers or DNS servers.

**Q7. How many SOA record does each zone contain?**

Each zone will have one SOA record. This records contains many miscellaneous settings for the zone, such as who is responsible for the zone, refresh interval settings, TTL (Time To Live) settings, and a serial number (incremented with every update).

**Q8. Short summary of the records in DNS.**

The NS records are used to point to additional DNS servers. The PTR record is used for reverse lookups (IP to name). CNAME records are used to give a host multiple names. MX records are used when configuring a domain for email.

**Q9. What is an AD-integrated zone?**

AD-integrated zones store the zone data in Active Directory and use the same replication process used to replicate other data between domain controllers. The one catch with AD-integrated zones is that the DNS server must also be a domain controller. Overloading DNS server responsibilities on your domain controllers may not be something you want to do if you plan on supporting a large volume of DNS requests.

**Q10.What is a STUB zone?**

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.

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.

**Q11. What does a stub zone consists of?**

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.

**Q12. How the resolution in a stub zone takes place?**

When a DNS client performs a recursive query operation on a DNS server hosting a stub zone, the DNS server uses the resource records in the stub zone to resolve the query. The DNS server sends an iterative query to the authoritative DNS servers specified in the NS resource records of the stub zone as if it were using NS resource records in its cache. If the DNS server cannot find the authoritative DNS servers in its stub zone, the DNS server hosting the stub zone attempts standard recursion using its root hints.

The DNS server will store the resource records it receives from the authoritative DNS servers listed in a stub zone in its cache, but it will not store these resource records in the stub zone itself; only the SOA, NS, and glue A resource records returned in response to the query are stored in the stub zone. The resource records stored in the cache are cached according to the Time-to-Live (TTL) value in each resource record. The SOA, NS, and glue A resource records, which are not written to cache, expire according to the expire interval specified in the stub zone's SOA record, which is created during the creation of the stub zone and updated during transfers to the stub zone from the original, primary zone.

If the query was an iterative query, the DNS server returns a referral containing the servers specified in the stub zone.

**Q 13.What is the benefits of Active Directory Integration?**

For networks deploying DNS to support Active Directory, directory-integrated primary zones are strongly recommended and provide the following benefits:

*** Multimaster update and enhanced security based on the capabilities of Active Directory**

In a standard zone storage model, DNS updates are conducted based upon a single-master update model. In this model, a single authoritative DNS server for a zone is designated as the primary source for the zone.

This server maintains the master copy of the zone in a local file. With this model, the primary server for the zone represents a single fixed point of failure. If this server is not available, update requests from DNS clients are not processed for the zone.

With directory-integrated storage, dynamic updates to DNS are conducted based upon a multimaster update model.

In this model, any authoritative DNS server, such as a domain controller running a DNS server, is designated as a primary source for the zone. Because the master copy of the zone is maintained in the Active Directory database, which is fully replicated to all domain controllers, the zone can be updated by the DNS servers operating at any domain controller for the domain.

With the multimaster update model of Active Directory, any of the primary servers for the directory-integrated zone can process requests from DNS clients to update the zone as long as a domain controller is available and reachable on the network.

Also, when using directory-integrated zones, you can use access control list (ACL) editing to secure a dnsZone object container in the directory tree. This feature provides granulated access to either the zone or a specified RR in the zone.

For example, an ACL for a zone RR can be restricted so that dynamic updates are only allowed for a specified client computer or a secure group such as a domain administrators group. This security feature is not available with standard primary zones.

Note that when you change the zone type to be directory-integrated, the default for updating the zone changes to allow only secure updates. Also, while you may use ACLs on DNS-related Active Directory objects, ACLs may only be applied to the DNS client service.

*** Directory replication is faster and more efficient than standard DNS replication.**

Because Active Directory replication processing is performed on a per-property basis, only relevant changes are propagated. This allows less data to be used and submitted in updates for directory-stored zones.

**Note**: Only primary zones can be stored in the directory. A DNS server cannot store secondary zones in the directory. It must store them in standard text files. The multimaster replication model of Active Directory removes the need for secondary zones when all zones are stored in Active Directory.

**Q14. What is Scavenging?**

DNS scavenging is the process whereby resource records are automatically removed if they are not updated after a period of time. Typically, this applies to only resource records that were added via DDNS, but you can also scavenge manually added, also referred to as static, records. DNS scavenging is a recommended practice so that your DNS zones are automatically kept clean of stale resource records.

**Q15. What is the default interval when DNS server will kick off the scavenging process?**

The default value is 168 hours, which is equivalent to 7 days.

**Q8. What's the largest number I can use in an MX record?**

> Could you tell us the highest possible number we can use for the MX
> preference?

Preference is an unsigned, 16-bit number, so the largest number you
can use is 65535.

**Q9. Why are there only 13 root name servers?**

> I'm very wondering why there are only 13 root servers on globally.
> Some documents explain that one of the reason is technical limit on Domain
> Name System (without any detailed explanation).
> From my understanding, it seems that some limitation of NS record numbers
> in DNS packet that specified by certain RFCs, or just Internet policy stuff.
> Which one is proper reason?

It's a technical limitation. UDP-based DNS messages can be up to 512 bytes
long, and only 13 NS records and their corresponding A records will fit into a DNS message that size.

**IMP information**
**http://www.menandmice.com/online_docs_and_faq/glossary/glossarytoc.htm**

10)Active directory Integrated DNS and non –integrated (In integrated DNS we can't edit DNS if failed need to restore System state backup and non Int DNS we can edit DNS and DNS files)

**11)What is recursion Authoritative DNS server?**

Ans ---

A DNS server which *recursive DNS servers* contact in order to *resolve* a given *DNS node*

DNS node

A name which DNS usually converts in to an IP, such as www.yahoo.com. Not all DNS nodes have IPs, however.

DNS record

A single piece of DNS data, which can either, be data for a DNS node, or meta-data which DNS uses.

DNS server

A program which resolves DNS records

DNS server administrator

A person who manages DNS; setting up DNS servers, changing DNS records, and what not.

Domain registry

A domain registry is a company that allows one to have their *authoritative DNS servers* be contacted by *recursive name servers*.

Domain suffix

The part of the domain which is (usually) after the first dot in a DNS node. The domain suffix for www.yahoo.com, for example, is yahoo.com.

Domain zone

A domain zone is a set of one or more DNS nodes. All names in a given domain zone share the same *domain suffix*.

IP

A number which a computer connected to the internet has, similar to a phone number.

Internet service provider

An internet service provider (or ISP) is a company that offers access to the internet.

Mail Transport Agent

A computer program which accepts incoming SMTP (email) connections, allowing a server to receive email.

Recursive DNS server

A recursive DNS server is a DNS server which contacts other DNS servers to *resolve* a given *DNS node*.

To resolve.

To convert a *DNS node*, such as www.yahoo.com, in to an IP, such as 10.17.243.32.

To serve

The action of an authoritative DNS server making DNS nodes available to recursive DNS servers.

Static IP address

A static IP address is an IP addresses whose value does not change. Only some internet service providers offer static IP addresses.

**12) Iterative DNS Resolution**

When a client sends an iterative request to a name server, the server responds back with either the answer to the request (for a regular resolution, the IP address we want) ***or*** the name of another server that has the information or is closer to it. The original client must then *iterate* by sending a new request to this referred server, which again may either answer it or provide another server name. The process continues until the right server is found; the method is illustrated in

**Iterative and recursion option is at DNS properties we can enable the same which we want (by default Recursion enabled)**

**If u enable Iterative then we need proxy server for Internet otherwise we will not able get Internet or the DNS request will not pass to internet from u r local DNS.**

**So that time proxy will take care of to connect Internet**

13) DNS round robin Is there at DNS properties we can enable that DNS query will fire like round robin (If we have two DNS server a client request take one by one we can say load balancing as well)

**Q3. What if a FSMO server fails?**

| | |
|---|---|
| Schema Master | No updates to the Active Directory schema will be possible. Since schema updates are rare (usually done by certain applications and possibly an Administrator adding an attribute to an object), then the malfunction of the server holding the Schema Master role will not pose a critical problem. |
| Domain Naming Master | The Domain Naming Master must be available when adding or removing a domain from the forest (i.e. running DCPROMO). If it is not, then the domain cannot be added or removed. It is also needed when promoting or demoting a server to/from a Domain Controller. Like the Schema Master, this functionality is only used on occasion and is not critical unless you are modifying your domain or forest structure. |
| PDC Emulator | The server holding the PDC emulator role will cause the most problems if it is unavailable. This would be most noticeable in a mixed mode domain where you are still running NT 4 BDCs and if you are using down-level clients (NT and Win9x). Since the PDC emulator acts as a NT 4 PDC, then any actions that depend on the PDC would be affected (User Manager for Domains, Server Manager, changing passwords, browsing and BDC replication).<br>In a native mode domain the failure of the PDC emulator isn't as critical because other domain controllers can assume most of the responsibilities of the PDC emulator. |
| RID Master | The RID Master provides RIDs for security principles (users, groups, computer accounts). The failure of this FSMO server would have little impact unless you are adding a very large number of users or groups. Each DC in the domain has a pool of RIDs already, and a problem would occur only if the DC you adding the users/groups on ran out of RIDs. |
| Infrastructure Master | This FSMO server is only relevant in a multi-domain environment. If you only have one domain, then the Infrastructure Master is irrelevant. Failure of this server in a multi-domain environment would be a problem if you are trying to add objects from one domain to another. |

**Q5. Can you Move FSMO roles?**

Yes, moving a FSMO server role is a manual process, it does not happen automatically. But what if you only have one domain controller in your domain? That is fine. If you have only one domain controller in your organization then you have one forest, one domain, and of course the one domain controller. All 5 FSMO server roles will exist on that DC. There is no rule that says you have to have one server for each FSMO server role.

**Q6. Where to place the FSMO roles?**

Assuming you do have multiple domain controllers in your domain, there are some best practices to follow for placing FSMO server roles.

- The Schema Master and Domain Naming Master should reside on the same server, and that machine should be a Global Catalog server. Since all three are, by default, on the first domain controller installed in a forest, then you can leave them as they are.

**Note:** According to MS, the Domain Naming master needs to be on a Global Catalog Server. If you are going to separate the Domain naming master and Schema master, just make sure they are both on Global Catalog servers.

**IMP:- Why Infrastructure Master should not be on the same server that acts as a Global Catalog server?**

- The Infrastructure Master should not be on the same server that acts as a Global Catalog server. The reason for this is the Global Catalog contains information about every object in the forest. When the Infrastructure Master, which is responsible for updating Active Directory information about cross domain object changes, needs information about objects not in it's domain, it contacts the Global Catalog server for this information. If they both reside on the same server, then the Infrastructure Master will never think there are changes to objects that reside in other domains because the Global Catalog will keep it constantly updated. This would result in the Infrastructure Master never replicating changes to other domain controllers in its domain.
**Note:** In a single domain environment this is not an issue.

- Microsoft also recommends that the PDC Emulator and RID Master be on the same server. This is not mandatory like the Infrastructure Master and the Global Catalog server above, but is recommended. Also, since the PDC Emulator will receive more traffic than any other FSMO role holder, it should be on a server that can handle the load.

- It is also recommended that all FSMO role holders be direct replication partners and they have high bandwidth connections to one another as well as a Global Catalog server.

**Q7.What permissions you should have in order to transfer a FSMO role?**

Before you can transfer a role, you must have the appropriate permissions depending on which role you plan to transfer:

| | |
|---|---|
| Schema Master | member of the Schema Admins group |
| Domain Naming Master | member of the Enterprise Admins group |
| PDC Emulator | member of the Domain Admins group and/or the Enterprise Admins group |
| RID Master | member of the Domain Admins group and/or the Enterprise Admins group |
| Infrastructure Master | member of the Domain Admins group and/or the Enterprise Admins group |

**FSMO TOOLS**

**Q8. Tools to find out what servers in your domain/forest hold what server roles?**

**1. Active Directory Users and Computers:-** use this snap-in to find out where the domain level FSMO roles are located (PDC Emulator, RID Master, Infrastructure Master), and also to change the location of one or more of these 3 FSMO roles.

Open Active Directory Users and Computers, right click on the domain you want to view the FSMO roles for and click "Operations Masters". A dialog box (below) will open with three tabs, one for each FSMO role. Click each tab to see what server that role resides on. To **change** the server roles, you must first connect to the domain controller you want to move it to. Do this by right clicking "Active Directory Users and Computers" at the top of the Active Directory Users and Computers snap-in and choose "Connect to Domain Controller". Once connected to the DC, go back into the Operations Masters dialog box, choose a role to move and click the Change button. When you do connect to another DC, you will notice the name of that DC will be in the field below the Change button (not in this graphic).

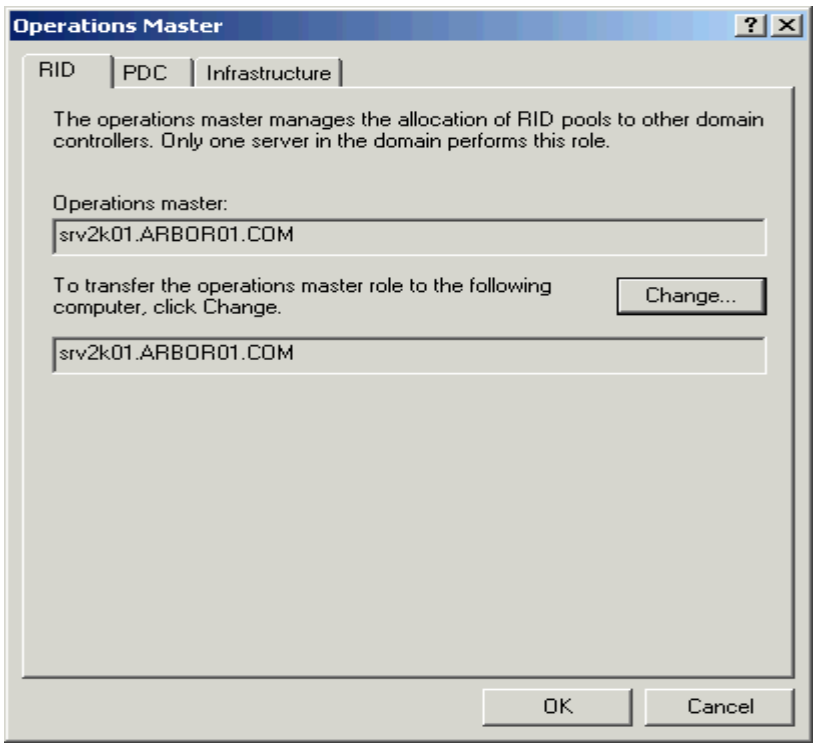


**2. Active Directory Domains and Trusts** - use this snap-in to find out where the Domain Naming Master FSMO role is and to change its location.

The process is the same as it is when viewing and changing the Domain level FSMO roles in Active Directory Users and Computers, except you use the Active Directory Domains and Trusts snap-in. Open Active Directory Domains and Trusts, right click "Active Directory Domains and Trusts" at the top of the tree, and choose "Operations Master". When you do, you will see the dialog box below. **Changing** the server that houses the Domain Naming Master requires that you first connect to the new domain controller, then click the Change button. You can connect to another domain controller by right clicking "Active Directory Domains and Trusts" at the top of the Active Directory Domains and Trusts snap-in and choosing "Connect to Domain Controller".

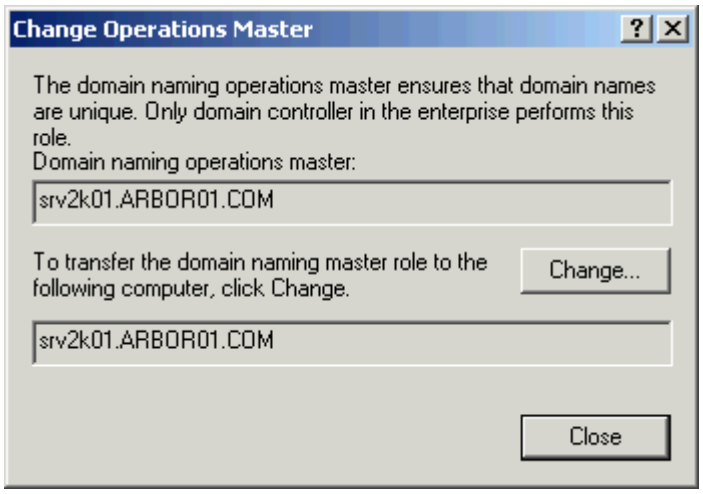


**3. Active Directory Schema** - this snap-in is used to view and change the Schema Master FSMO role. However... the Active Directory Schema snap-in is not part of the default Windows 2000 administrative tools or installation. You first have to install the Support Tools from the \Support directory on the Windows 2000 server CD or install the Windows 2000 Server Resource Kit. Once you install the support tools you can open up a blank Microsoft Management Console (start, run, mmc) and add the snap-in to the console. Once the snap-in is open, right click "Active Directory

Schema" at the top of the tree and choose "Operations Masters". You will see the dialog box below. **Changing** the server the Schema Master resides on requires you first connect to another domain controller, and then click the Change button.

You can connect to another domain controller by right clicking "Active Directory Schema" at the top of the Active Directory Schema snap-in and choosing "Connect to Domain Controller".

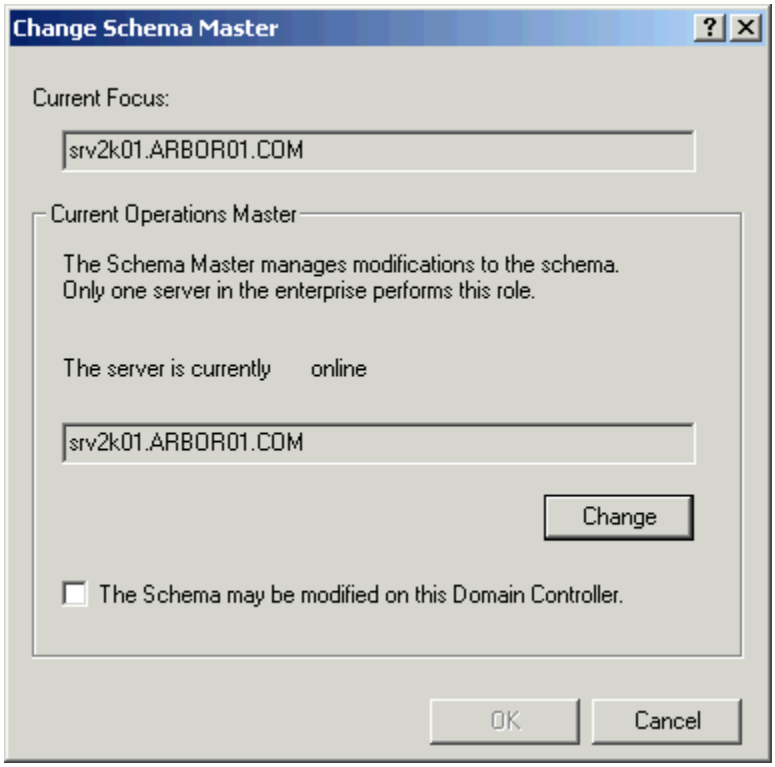


**4. Netdom**

The easiest and fastest way to find out what server holds what FSMO role is by using the **Netdom** command line utility. Like the Active Directory Schema snap-in, the Netdom utility is only available if you have installed the Support Tools from the Windows 2000 CD or the Win2K Server Resource Kit.

To use Netcom to view the FSMO role holders, open a command prompt window and type: netdom query fsmo and press enter. You will see a list of the FSMO role servers:

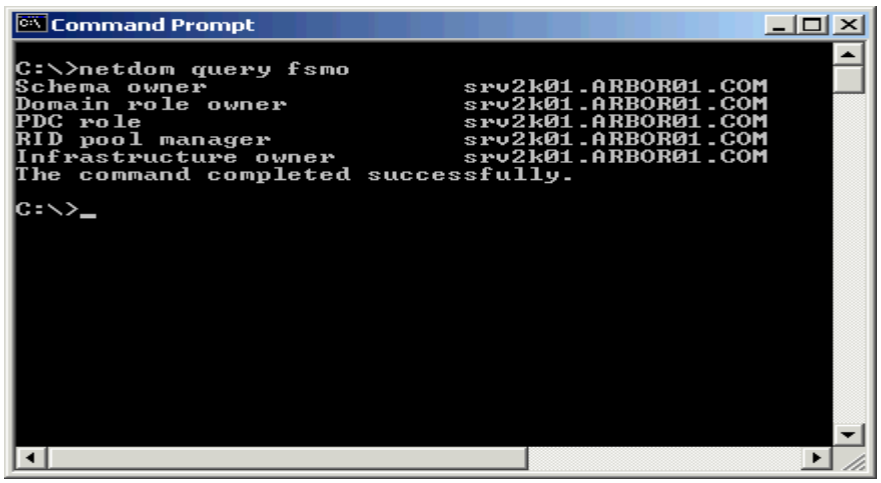

**5. Active Directory Relication Monitor**

Another tool that comes with the Support Tools is the **Active Directory Relication Monitor**. Open this utility from Start, Programs, Windows 2000 Support Tools. Once open, click Edit, Add Monitored Server and add the name of a Domain Controller. Once added right click the Server name and choose properties. Click the FSMO Roles tab to view the servers holding the 5 FSMO roles (below). You cannot change roles using Replication Monitor, but this tool has many other useful purposes in regard to Active Directory information. It is something you should check out if you haven't already.

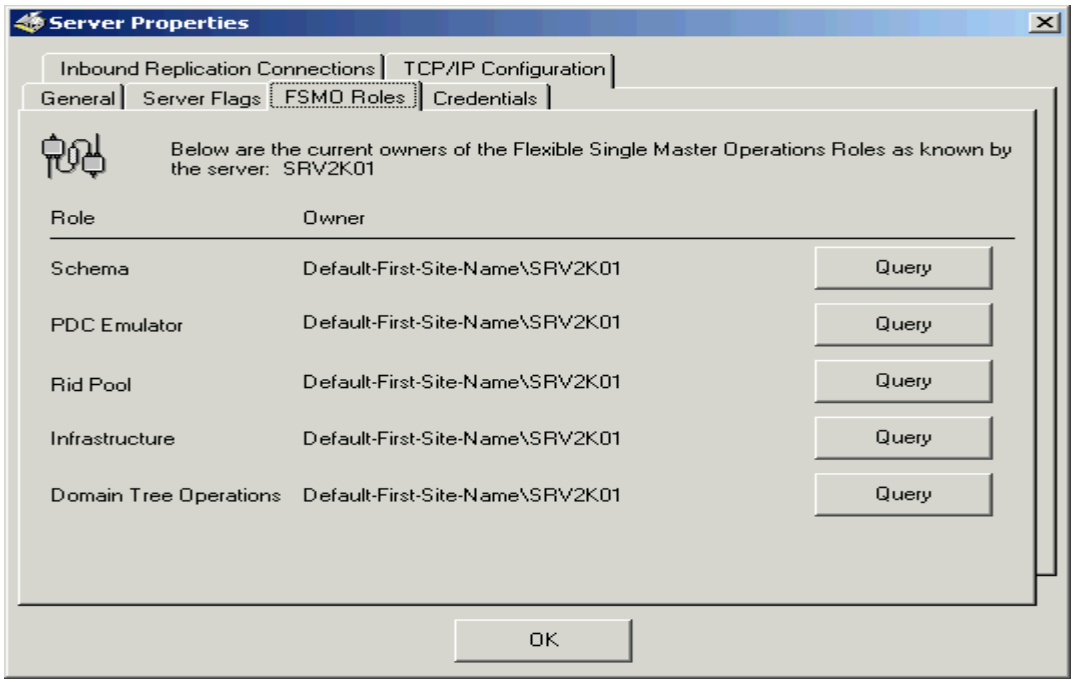


Finally, you can use the **Ntdsutil.exe utility** to gather information about and change servers for FSMO roles. Ntdsutil.exe, a command line utility that is installed with Windows 2000 server, is rather complicated and beyond the scope of this document.

**6. DUMPFSMOS**

Command-line tool to query for the current FSMO role holders

Part of the Microsoft Windows 2000 Server Resource Kit

Downloadable from http://www.microsoft.com/windows2000

/techinfo/reskit/default.asp

Prints to the screen, the current FSMO holders

Calls NTDSUTIL to get this information

**7. NLTEST**

Command-line tool to perform common network administrative tasks

Type “nltest /?” for syntax and switches

Common uses

Get a list of all DCs in the domain

Get the name of the PDC emulator

Query or reset the secure channel for a server

Call DsGetDCName to query for an available domain controller

**8. Adcheck (470k)** (3rd party)

A simple utility to view information about AD and FSMO roles

http://www.svrops.com/svrops/downloads/zipfiles/ADcheck.msi

**Q1. What are Group Policies?**

Group Policies are settings that can be applied to Windows computers, users or both. In Windows 2000 there are hundreds of Group Policy settings. Group Policies are usually used to lock down some aspect of a PC. Whether you don't want users to run Windows Update or change their Display Settings, or you want to insure certain applications are installed on computers - all this can be done with Group Policies.

Group Policies can be configured either **Locally** or by **Domain Polices**. Local policies can be accessed by clicking Start, Run and typing gpedit.msc. They can also be accessed by opening the Microsoft Management Console (Start, Run type mmc), and adding the Group Policy snap-in. You must be an Administrator to configure/modify Group Policies. Windows 2000 Group Policies can only be used on Windows 2000 computers or Windows XP computers. They cannot be used on Win9x or WinNT computers.

**Q2. Domain policy gets applied to whom?**

Domain Policies are applied to computers and users who are members of a Domain, and these policies are configured on **Domain Controllers**. You can access Domain Group Polices by opening Active Directory Sites and Services (these policies apply to the Site level only) or Active Directory Users and Computers (these policies apply to the Domain and/or Organizational Units).

**Q3. From Where to create a Group Policy?**

To create a Domain Group Policy Object open Active Directory Sites and Services and right click Default-First-Site-Name or another Site name, choose properties, then the Group Policy tab, then click the **New button**. Give the the GPO a name, then click the **Edit button** to configure the policies.
For Active Directory Users and Computers, it the same process except you right click the Domain or an OU and choose properties.

**Q4. Who can Create/Modify Group Policies?**

You have to have Administrative privileges to create/modify group policies. The following table shows who can create/modify group policies:

| Policy Type | Allowable Groups/Users |
|---|---|
| Site Level Group Policies | Enterprise Administrators and/or Domain Administrators in the root domain. The root domain is the first domain created in a tree or forest. The Enterprise Administrators group is found only in the root domain. |
| Domain Level Group Policies | Enterprise Administrators, Domain Administrators or members of the built-in group - Group Policy Creator Owners. By default only the Administrator user account is a member of this group |
| OU Level Group Policies | Enterprise Administrators, Domain Administrators or members of the Group Policy Creator Owners. By default only the Administrator user account is a member of this group.<br><br>Additionally, at the OU level, users can be delegated control for the OU Group Policies by starting the Delegate Control |

| | |
|---|---|
| | Wizard (right click the OU and choose Delegate Control). However, the wizard only allows the delegated user to Link already created group policies to the OU. If you want to give the OU administrators control over creating/modifying group policies, add them to the Group Policy Creator Owners group for the domain. |
| Local Group Policies | The local Administrator user account or members of the local Administrators group. |

**Q5. How are Group Policies Applied?**

Group Polices can be configured locally, at the Site level, the Domain level or at the Organizational Unit (OU) level. Group Policies are applied in a Specific Order, LSDO - **Local** policies first, then **Site** based policies, then **Domain** level policies, then **OU** polices, then **nested OU** polices (OUs within OUs). Group polices cannot be linked to a specific user or group, only container objects.

In order to apply Group Polices to specific users or computers, you add users (or groups) and computers to container objects. Anything in the container object will then get the policies linked to that container. Sites, Domains and OUs are considered container objects.

Computer and User Active Directory objects **do not** have to put in the same container object. For example, Sally the user is an object in Active Directory. Sally's Windows 2000 Pro PC is also an object in Active Directory. Sally the user object can be in one OU, while her computer object can be another OU. It all depends on how you organize your Active Directory structure and what Group Policies you want applied to what objects.

**User and Computer Policies**

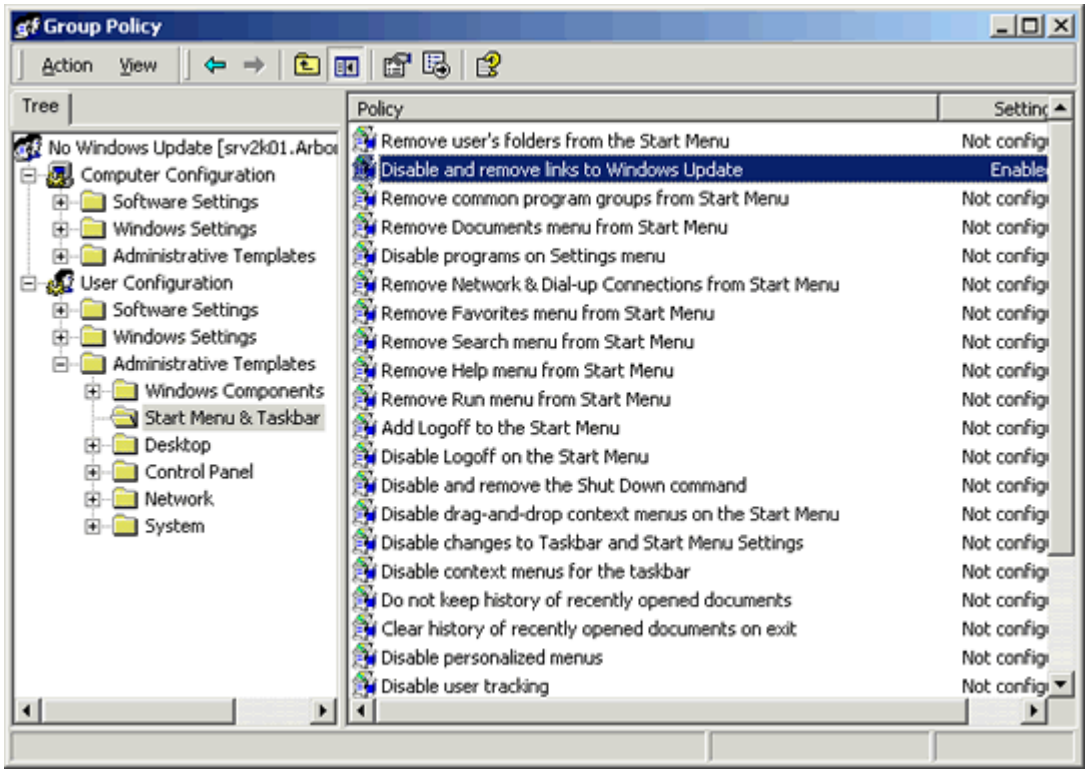


There are two nodes in each Group Policy Object that is created. A **Computer** node and a **User** Node. They are called **Computer Configuration** and **User Configuration** (see image above). The polices configured in the Computer node apply to the computer as a whole. Whoever logs onto that computer will see those policies.

**Note:** Computer policies are also referred to as machine policies.

User policies are user specific. They only apply to the user that is logged on. When creating Domain Group Polices you can disable either the Computer node or User node of the Group Policy Object you are creating. By disabling a node that no policies are defined for, you are decreasing the time it takes to apply the polices.
**To disable the node polices:** After creating a Group Policy Object, click that Group Policy Object on the Group Policy tab, then click the Properties button. You will see two check boxes at the bottom of the General tab.

It's important to understand that when Group Policies are being applied, **all the policies** for a node are evaluated first, and then applied. They are not applied one after the other. For example, say Sally the user is a member of the Development OU, and the Security OU. When Sally logs onto her PC the policies set in the User node of the both the Development OU and the Security OU Group Policy Objects are evaluated, as a whole, and then applied to Sally the user. They are not applied Development OU first, and then Security OU (or visa- versa).
The same goes for Computer policies. When a computer boots up, all the Computer node polices for that computer are evaluated, then applied.

When computers **boot up**, the Computer policies are applied. When users **login**, the User policies are applied. When user and computer group policies overlap, the **computer policy wins**.

**Note:** IPSec and EFS policies are not additive. The last policy applied is the policy the user/computer will have.

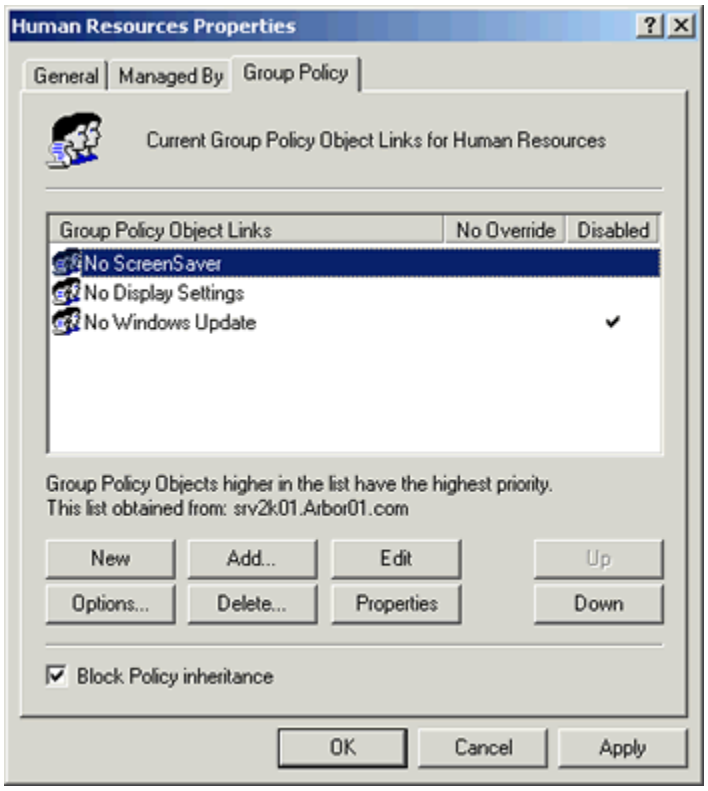


When applying multiple Group Policies Objects from any container, Group Policies are applied from bottom to top in the Group Policy Object list. The top Group Policy in the list is the last to be applied. In the above image you can see three Group Policy Objects associated with the Human Resources OU. These polices would be applied No Windows Update first, then No Display Settings, then No Screensaver. If there were any conflicts in the policy settings, the one above it would take precedence.

**Q6.How to disable Group Policy Objects**

When you are creating a Group Policy Object, the changes happen immediately. There is no "saving" of GPOs. To prevent a partial GPO from being applied, **disable the GPO** while you are configuring it. To do this, click the Group Policy Object on the Group Policy tab and under the Disable column, double click - a little check will appear. Click the Edit button, make your changes, then double click under the Disable column to re-enable the GPO. Also, if you want to temporarily disable a GPO for troubleshooting reasons, this is the place to do it. You can also click the Options button on the Group Policy tab and select the Disabled check box.

**Q7. When does the group policy Scripts run?**

**Startup** scripts are processed at computer bootup and before the user logs in.
**Shutdown** scripts are processed after a user logs off, but before the computer shuts down.

**Login** scripts are processed when the user logs in.
**Logoff** scripts are processed when the user logs off, but before the shutdown script runs.

**Q8. When the group policy gets refreshed/applied?**

Group Policies can be applied when a computer boots up, and/or when a user logs in. However, policies are also refreshed automatically according to a predefined schedule. This is called **Background Refresh**.

**Background** refresh for **non DCs** (PCs and Member Servers) is every 90 mins., with a +/- 30 min. interval. So the refresh could be 60, 90 or 120 mins. For **DCs** (Domain Controllers), background refresh is every **5 mins**.
Also, every **16 hours** every PC will request all group policies to be reapplied (user and machine)
These settings can be changed under Computer and User Nodes, Administrative Templates, System, Group Policy.

**Q9. Which are the policies which do not get affected by background refresh?**

Policies not affected by background refresh. These policies are only applied at **logon time**:

Folder Redirection
Software Installation
Logon, Logoff, Startup, Shutdown Scripts

**Q10. How to refresh Group Policies suing the command line?**

Secedit.exe is a command line tool that can be used to refresh group policies on a Windows 2000 computer. To use secedit, open a command prompt and type:

***secedit /refreshpolicy user policy*** to refresh the user policies
***secedit /refreshpolicy machine_policy*** to refresh the machine (or computer) policies

These parameters will only refresh any user or computer policies that have changed since the last refresh. To force a reload of all group policies regardless of the last change, use:

***secedit /refreshpolicy user_policy /enforce***
***secedit /refreshpolicy machine_policy /enforce***

Gpupdate.exe is a command line tool that can be used to refresh group policies on a Windows XP computer. It has replaced the secedit command. To use gpupdate, open a command prompt and type:

***gpupdate /target:user*** to refresh the user policies
***gpupdate /target:machine*** to refresh the machine (or computer) policies

As with secedit, these parameters will only refresh any user or computer policies that have changed since the last refresh. To force a reload of all group policies regardless of the last change, use:

*gpupdate /force*

Notice the /force switch applies to both user and computer policies. There is no separation of the two like there is with secedit

**Q11. What is the Default Setting for Dial-up users?**

Win2000 considers a slow dial-up link as anything less than 500kbps. When a user logs into a domain on a link under 500k some policies are not applied.

Windows 2000 will automatically detect the speed of the dial-up connection and make a decision about applying Group Policies.

**Q12. Which are the policies which get applied regardless of the speed of the dial-up connection?**

Some policies are always applied regardless of the speed of the dial-up connection. These are:

Administrative Templates
Security Settings
EFS Recovery
IPSec

**Q13. Which are the policies which do not get applied over slow links?**

IE Maintenance Settings
Folder Redirection
Scripts
Disk Quota settings
Software Installation and Maintenance

These settings can be changed under Computer and User Nodes, Administrative Templates, System, Group Policy.

If the user connects to the domain using "Logon Using Dial-up Connection" from the logon screen, once the user is authenticated, the computer policies are applied first, followed by the user policies.

If the user connects to the domain using "Network and Dial-up Connections", **after they logon**, the policies are applied using the standard refresh cycle.

**Q14. Which are the two types of default policies?**

There are **two default** group policy objects that are created when a domain is created. The Default Domain policy and the Default Domain Controllers policy.

**Default Domain Policy** - this GPO can be found under the group policy tab for that domain. It is the first policy listed. The default domain policy is unique in that certain policies can only be applied at the domain level.

If you double click this GPO and drill down to Computer Configuration, Windows Settings, Security Settings, Account Policies, you will see three policies listed:

Password Policy
Acount Lockout Policy
Kerberos Policy

These 3 policies can only be set at the domain level. If you set these policies anywhere else- Site or OU, they are ignored. **However**, setting these 3 policies at the OU level will have the effect of setting these policies for users who log on *locally* to their PCs. Login to the domain you get the domain policy, login locally you get the OU policy.

If you drill down to Computer Configuration, Windows Settings, Security Settings, Local Policies, Security Options, there are 3 policies that are affected by Default Domain Policy:

Automatically log off users when logon time expires

Rename Adminsitrator Account - When set at the domain level, it affects the Domain Administrator account only.
Rename Guest Account - When set at the domain level, it affects the Domain Guest account only.

The Default Domain Policy should be used only for the policies listed above. If you want to create additional domain level policies, you should create additional domain level GPOs.
Do not delete the Default Domain Policy. You can disable it, but it is not recommended.

**Defaults Domain Controllers Policy** - This policy can be found by right clicking the Domain Controllers OU, choosing Properties, then the Group Policy tab. This policy affects all Domain Controllers in the domain regardless of where you put the domain controllers. That is, no matter where you put your domain controllers in Active Directory (whatever OU you put them in), they will still process this policy.

Use the Default Domain Controllers Policy to set **local policies** for your domain controllers, e.g. Audit Policies, Event Log settings, who can logon locally and so on.

**Q15.How to restore Group policy setting back to default?**

The following command would replace both the Default Domain Security Policy and Default Domain Controller Security Policy. You can specify Domain or DC instead of Both, to only restore one or the other.

> dcgpofix /target:Both

**Note** that this must be run from a domain controller in the target domain where you want to reset the GPO

If you've ever made changes to the default GPOs and would like to revert back to the original settings, the dcgpofix utility is your solution. dcgpofix works with a particular version of schema. If the version it expects to be current is different from what is in Active Directory, it not restore the GPOs. You can work around this by using the /ignoreschema switch, which restore the GPO according to the version dcgpofix thinks is current. The only time you might experience this issue is if you install a service pack on a domain controller (dc1) that extends schema, but have not installed it yet on a second domain controller (dc2). If you try to run

dcgpofix from dc2, you will receive the error since a new version of the schema and the dcgpofix utility was installed on dc1.

18)Where store DHCP database and what is file name –I think DHCPS.mdb --Pls ckeck this
How to configure DHCP for Command promotes "net stat" command can use …….

**Kerberos v.5**

Windows 2000 Active Directory relies on a different authentication protocol than Windows NT 4. Where NT 4 used the NT Lan Manager (NTLM) protocol for authentication, Windows 2000 utilizes Kerberos. The Kerberos protocol was developed at MIT, and is named after Cerebus, the three-headed fire-breathing dog that guards the gate to Hades. Why do I bother telling you this? Because it makes it easier to remember that Kerberos is a 3-pronged authentication scheme. The three parts of Kerberos are:

1.Client - the system/user making the request
2.Server - the system that offers a service to systems whose identity can be confirmed
3.Key Distribution Center (KDC) - the third-party intermediary between the client and the server, who vouches for the identity of a client. In a Windows 2000 environment, the KDC in a domain controller running Active Directory (It could be a UNIX-based KDC also)

The way that Kerberos works can seem a little intimidating if you get into all of the tiny little details, but I'll spare them for an overview of how things work. It is more important that you understand the process to begin with. If you want every behind-the-scenes detail, I've provided a link at the end of the section.

In a Kerberos environment a user provides a username, password, and domain name (often referred to as a Realm in Kerberos lingo) that they wish to log on to. This information is sent to a KDC, who authenticate the user. If the user is valid, they are presented with something called a ticket-granting ticket, or TGT. I like to consider the TGT to be like a hand-stamp admission to a country fair - it proves that you have paid admission and have proof. The TGT is helpful in that it does not require you to constantly re-authenticate every time you need to access a server.

However, if you do want to access a server, you still require a ticket for that server or you will not be able to create a session with that machine. Think of a ticket as being like the ticket you need to purchase to get on rides at the country fair - even though you've paid admission (proved by the hand-stamp or TGT), you still need a ticket to get into the haunted house. When you wish to access a server, you first need to go to the KDC, present your TGT as proof of identity, and then request a session ticket for the server you wish to contact. This ticket simply acts as authentication between the client and the server you wish to contact. If you are authenticated, whether or not you will be able to actually access anything on the server will depend on your permissions. The TGT and session tickets that you are presented with actually expire after a period of time that is configurable via group policy. The default value for a TGT (also referred to as a user ticket) is 7 days, while the default value for a session ticket (sometimes called a service ticket) is 10 hours.

In a single-domain environment, Kerberos authentication is pretty straightforward. However in a multiple domain environment Kerberos has more steps involved. The reason for this is that when you are attempting to obtain a session ticket for a server, it must be obtained from a KDC in the domain where the server exists. Also, you must obtain session tickets in order to traverse the trust-path to the KDC you need to contact. The example below outlines the steps necessary for a client in west.win2000trainer.com to access a server in east.win2000trainer.com.

1.The client logs on to the network as a user in east.win2000trainer.com, and is presented with a TGT.
2.The client wants to communicate with a server in west.win2000trainer.com. It contacts the KDC in east.win2000trainer.com, asking for a session ticket for a KDC in the win2000trainer.com domain.
3. After it receives this ticket, it contacts the KDC in win2000trainer.com, requesting a session ticket for the KDC in west.win2000trainer.com.
4. After it receives this ticket, it contacts the KDC in west.win2000trainer.com, and requests a session ticket for the server in west.win2000trainer.com whom it originally wanted to contact.
5.Once granted the session ticket for the server, the client contacts that server directly and can access resources according to the permissions in place.

If this seems like a great deal of steps, that is indeed true. This is one of the reasons that you might consider implementing shortcut trusts, as outlined in my last article. If shortcut trusts exist, the shorter available path would be used. Kerberos is a wonderful protocol in that is makes the network much more secure, due to the necessity of authentication between clients and servers before a session can be established. It is actually much faster than you might think. For a good hands-on experiment, you might consider setting up multiple domains and then running network monitor while accessing resources between domains. Though the packets contents are encrypted, it will still give you a great idea of what is happening behind the scenes. Three utilities that you should be aware of for troubleshooting Kerberos problems are Netdom (discussed in a previous article), as well as the resource kit utilities KerbTray.exe and Llist.exe.

| | NTLM | Kerberos |
|---|---|---|
| Cryptographic technology | Symmetric cryptography | Basic Kerberos: symmetric cryptography<br>Kerberos PKINIT: symmetric and asymmetric cryptography |
| Trusted third party | Domain controller | Basic Kerberos: domain controller with KDC service<br>Kerberos PKINIT: domain controller with KDC service and Enterprise CA |
| Microsoft-supported platforms | Windows 95, Windows 98, Windows ME, Windows NT4, Windows 2000, Windows XP, and Windows Server 2003 | Windows 2000, Windows XP, and Windows Server 2003* |
| Features | Slower authentication because of pass-through authentication | Faster authentication because of unique ticketing system |
| | No mutual authentication | Mutual authentication |
| | No support for delegation of authentication | Support for delegation of authentication |
| | No support for smart card logon feature | Support for smart card logon feature |
| | Proprietary Microsoft authentication protocol | Open standard |
| | No protection for authorization data carried in NTLM messages† | Cryptographic protection for authorization data carried in Kerberos tickets |

* Remember from the previous chapter that Kerberos can only be used for domain logon to a Windows 2000 or Windows Server 2003 domain.
† This was the case for NTLM version 1; this problem has been resolved in NTLM version 2.

**Active dir…integrated**

The DNS service provided with Windows 2000 Server meets both these requirements and also offers two important additional features:

- Active Directory integration

  Using this feature, the Windows 2000 DNS service stores zone data in the directory. This makes DNS replication create multiple masters, and it allows any DNS server to accept updates for a directory service-integrated zone. Using Active Directory integration also reduces the need to maintain a separate DNS zone transfer replication topology.

- Secure dynamic update

  Secure dynamic update is integrated with Windows security. It allows an administrator to precisely control which computers can update which names, and it prevents unauthorized computers from obtaining existing names from DNS.

The remaining DNS servers on your network that are not authoritative for the locator records do not need to meet these requirements. Servers that are not authoritative are generally able to answer SRV record queries even if they do not explicitly support that record type.

7.)**What are the RAID levels, Description? Min Max Discs supported by each level?**

RAID stands for *Redundant Array of Inexpensive Disks*. A RAID system consists of two or more disks working in parallel. They appear as one drive to the user, and offer enhanced performance or security (or both).
a) **RAID 0: striping b) RAID 1: mirroring c) RAID 3 d) RAID 5**
**e) RAID 10: a mix of RAID 0 & RAID 1**
a) RAID 0 :
In a RAID 0 system, data are split up in blocks that get written across all the drives in the array. By using multiple disks (at least 2) at the same time, RAID 0 offers superior I/O performance. This performance can be enhanced further by using multiple controllers, ideally one controller per disk.
**Advantages**

RAIDS 0 offers great performances, both in read and write operations. There is no overhead caused by parity controls.

All storage capacity can be used, there is no disk overhead.

The technology is easy to implement.

**Disadvantages**

RAID 0 is not fault-tolerant. If one disk fails, all data in the RAID 0 array are lost. It should not be used on mission-critical systems.

## Recommended Applications

- Video Production and Editing
- Image Editing
- Pre-Press Applications
- Any application requiring high bandwidth

B) RAID 1:

Data are stored twice by writing them to either the data disk (or set of data disks) and a mirror disk (or set of disks). If a disk fails, the controller uses either the data drive or the mirror drive for data recovery and continues operation. You need at least 2 disks for a RAID 1 array

RAID 1 systems are often combined with RAID 0 to improve performance. Such a system is sometimes referred to by the combined number: a RAID 10 system.

**Advantages**

RAID 1 offers excellent read speed and a write-speed that is comparable to that of a single disk.

In case a disk fails, data do not have to be rebuilding, they just have to be copied to the replacement disk.

RAID 1 is a very simple technology.

**Disadvantages**

The main disadvantage is that the effective storage capacity is only half of the total disk capacity because all data get written twice.

Software RAID 1 solution do not always allow a hot swap of a failed disk (meaning it cannot be replaced while the server keeps running). Ideally a hardware controller is used.

## Recommended Applications

- Accounting
- Payroll
- Financial
- Any application requiring very high availability

C) RAID 3:

On RAID 3 systems, data blocks are subdivided (striped) and written in parallel on two or more drives. An additional drive stores parity information. You need at least 3 disks for a RAID 3 array.

Since parity is used, a RAID 3 stripe set can withstand a single disk failure without losing data or access to data.

**Advantages**

RAID-3 provides high throughput (both read and write) for large data transfers.

Disk failures do not significantly slow down throughput.

**Disadvantages**

This technology is fairly complex and too resource intensive to be done in software.

lPerformance is slower for random, small I/O operations.

D) RAID 5:

RAID 5 is the most common secure RAID level. It is similar to RAID-3 except that data are transferred to disks by independent read and write operations (not in parallel). The data chunks that are written are also larger. Instead of a dedicated parity disk, parity information is spread across all the drives. You need at least 3 disks for a RAID 5 array.

A RAID 5 array can withstand a single disk failure without losing data or access to data. Although RAID 5 can be achieved in software, a hardware controller is recommended. Often extra cache memory is used on these controllers to improve the write performance

**Advantages**

Read data transactions are very fast while write data transaction are somewhat slower (due to the parity that has to be calculated).

**Disadvantages**

Disk failures have an effect on throughput, although this is still acceptable.

Like RAID 3, this is complex technology.

## Recommended Applications

- File and Application servers
- Database servers
- Web, E-mail, and News servers
- Intranet servers
- Most versatile RAID level

**E) RAID 10: a mix of RAID 0 & RAID 1**

RAID 10 combines the advantages (and disadvantages) of RAID 0 and RAID 1 in a single system. It provides security by mirroring all data on a secondary set of disks (disk 3 and 4 in the drawing below) while using striping across each set of disks to speed up datatransfers.

## Recommended Applications

- Database server requiring high performance and fault tolerance

**8.)What is Parity Bit?**

In computers, parity (from the Latin *paritas*: equal or equivalent) refers to a technique of checking whether data has been lost or written over when it's moved from one place in storage to another or when transmitted between computers

One can improve upon memory-style ECC disk arrays by noting that, unlike memory component failures, disk controllers can easily identify which disk has failed. Thus, one can use a single parity rather than a set of parity disks to recover lost information.

**21.) What are the partitions in AD ?**

Schema partition, configuration partition, global partition and Application Partition

**22.)What are the protocols in AD ?**

NTLM/KERBOROS V4,V5

**23.)What is a print server? Can we use AD as a print server ?**

Print servers probably show the greatest variation of machine, from dedicated print servers, you get printers hanging off domain controllers to 'Jet Direct' printers with their own network cards. In my experience there is a contrast between the software settings which are easy to configure and the hardware which constantly cries for attention e.g. paper jam, 'out of toner'.

**27.)What is the use of a OU in AD?**

As defined in the RFP for the LDAP standard, organizational units (OUs) are containers that logically store directory information and provide a method of addressing Active Directory through LDAP. In Active Directory, OUs are the primary method for organizing user, computer, and other object information into a more easily understandable layout. the organization has a root organizational unit where three nested organizational units (marketing, IT, and research) have been placed. This nesting enables the organization to distribute users across multiple containers for easier viewing and administration of network resources.

Organizational Units allow organizations on campus to maintain control over their resources, like computers, servers, printers and file shares. Below are instructions on how to maintain OUs.

## Via internet

## Active Directory Group Types

Two types of groups can be created in Active Directory. Each group type is used for a different purpose. A security group is one that is created for security purposes, while a distribution group is one created for purposes other than security purposes. Security groups are typically created to assign permissions, while distribution groups are usually created to distribute bulk e-mail to users. As one may notice, the main difference between the two groups is the manner in which each group type is used. Active Directory allows users to convert a security group into a distribution group and to convert a distribution group into a security group if the domain functional level is raised to Windows 2000 Native or above.

**Security Groups**

A security group is a collection of users who have the same permissions to resources and the same rights to perform certain system tasks. These are the groups to which permissions are assigned so that its members can access resources. Security groups therefore remove the need for an Administrator to individually assign permissions to users. Users that need to perform certain tasks can be grouped in a security group then assigned the necessary permissions to perform these tasks. Each user that is a member of the group has the same permissions. In addition to this, each group member receives any e-mail sent to a security group. When a security group is first created, it receives an SID. It is this SID that enables permissions to be assigned to security groups - the SID can be

included in a resource's DACL. An access token is created when a user logs on to the system. The access token contains the user's SID and the SID of those groups to which the user is a member of. This access token is referenced when the user attempts to access a resource. The access token is compared with the resource's DACL to determine which permissions the user should receive for the resource.

**Distribution Groups**

Distribution groups are created to share information with a group of users through e-mail messages. Thus, a distribution group is not created for security purposes. A distribution does not obtain an SID when it is created. Distribution groups enable the same message to be simultaneously sent to its group members. Messages do not need to be individually sent to each user. Applications such as Microsoft Exchange that work with Active Directory can use distribution groups to send bulk e-mail to groups of users.

**Dynamic Distribution Groups**

Dynamic distribution groups are mail-enabled Active Directory group objects that are created to expedite the mass sending of email messages and other information within a Microsoft Exchange organization. Unlike regular distribution groups that contain a defined set of members, the membership list for dynamic distribution groups is calculated each time a message is sent to the group, based on the filters and conditions that you define. When an email message is sent to a dynamic distribution group, it's delivered to all recipients in the organization that match the criteria defined for that group.

**Groups**

**Distribution Groups** — Used for email. Useful for programs such as MS Exchange.
**Security Groups** – Used to secure file/folders, printers, etc.

Local – Stored on the local SAM (Local Computers)
Domain Local – Stored on Domain Controllers.
Global Groups – Gives you a greater group scope.
Universal – Gives you an even broader group scope.

**Group Scopes**

Group scope normally describes the type of users that should be clubbed together in a way that is easy for their administration. Therefore, groups play an important part in domain. One group can be a member of other group(s), which is known as Group nesting. One or more groups can be members of any group in the entire domain(s) within a forest.

- **Domain Local Group**: Use this scope to grant permissions to domain resources that are located in the same domain in which the domain local group was created. Domain local groups can exist in all mixed, native, and interim functional level of domains and forests. Domain local group memberships are not limited as users can add members as user accounts and universal and global groups from any domain. Nesting

cannot be done in a domain local group. A domain local group will not be a member of another Domain Local or any other groups in the same domain.

- **Global Group**: Users with similar functions can be grouped under global scope and can be given permission to access a resource (like a printer or shared folder and files) available in local or another domain in the same forest. Simply put, global groups can be use to grant permissions to gain access to resources that are located in any domain but in a single forest as their memberships are limited. User accounts and global groups can be added only from the domain in which the global group is created. Nesting is possible in Global groups within other groups as users can add a global group into another global group from any domain. They can be members of a Domain Local group to provide permission to domain specific resources (like printers and published folder). Global groups exist in all mixed, native, and interim functional level of domains and forests.
- **Universal Group Scope**: these groups are precisely used for email distribution and can be granted access to resources in all trusted domain as these groups can only be used as a security principal (security group type) in a windows 2000 native or windows server 2003 domain functional level domain. Universal group memberships are not limited like global groups. All domain user accounts and groups can be a member of a universal group. Universal groups can be nested under a global or Domain Local group in any domain