DNS Zones Overview

A DNS zone is the contiguous portion of the DNS domain name space over which a DNS server has authority. A zone is a portion of a namespace. It is not a domain. A domain is a branch of the DNS namespace. A DNS zone can contain one or more contiguous domains. A DNS server can be authoritative for multiple DNS zones. A non-contiguous namespace cannot be a DNS zone.

A zone contains the resource records for all of the names within the particular zone. Zone files are used if DNS data is not integrated with Active Directory. The zone files contain the DNS database resource records that define the zone. If DNS and Active Directory are integrated, then DNS data is stored in Active Directory.

The different types of zones used in Windows Server 2003 DNS are listed below:

- Primary zone
- Secondary zone
- Active Directory-integrated zone
- Reverse lookup zone
- Stub zone

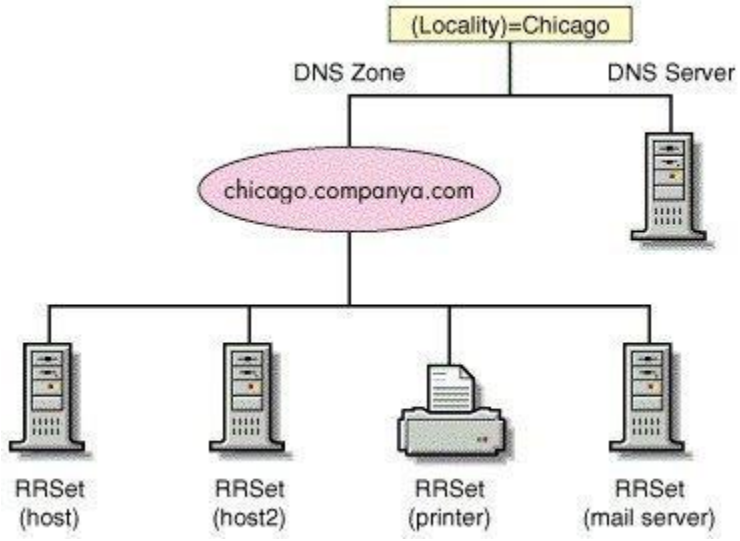

A *primary zone* is the only zone type that can be edited or updated because the data in the zone is the original source of the data for all domains in the zone. Updates made to the primary zone are made by the DNS server that is authoritative for the specific primary zone. Users can also back up data from a primary zone to a secondary zone.

A *secondary zone* is a read-only copy of the zone that was copied from the master server during zone transfer. In fact, a secondary zone can only be updated through zone transfer.

An *Active Directory-integrated zone* is a zone that stores its data in Active Directory. DNS zone files are not needed. This type of zone is an authoritative primary zone. An Active Directory-integrated zone's zone data is replicated during the Active Directory replication process. Active Directory-integrated zones also enjoy the Active Directory's security features.

A *reverse lookup* zone is an authoritative DNS zone. These zones mainly resolve IP addresses to resource names on the network. A reverse lookup zone can be either of the following zones:

- Primary zone
- Secondary zone

- Active Directory-integrated zone

A *stub zone* is a new Windows Server 2003 feature. Stub zones only contain those resource records necessary to identify the authoritative DNS servers for the master zone. Stub zones therefore contain only a copy of a zone, and are used to resolve recursive and iterative queries:

- *Iterative queries*: The DNS server provides the best answer it can. This can be:
  - The resolved name
  - A referral to a different DNS server
- *Recursive queries*: The DNS server has to reply with the requested information or with an error. The DNS server cannot provide a referral to a different DNS server.

Stub zones contain the following information:

- Start of Authority (SOA) resource records of the zone
- Resource records that list the authoritative DNS servers of the zone
- Glue address (A) resource records that are necessary for contacting the authoritative servers of the zone.

Zone delegation occurs when users assign authority over portions of the DNS namespace to subdomains of the DNS namespace. Users should delegate a zone under the following circumstances:

- To delegate administration of a DNS domain to a department or branch of the organization.
- To improve performance and fault tolerance of the DNS environment. Users can distribute DNS database management and maintenance between several DNS servers.

Understanding DNS Zone Transfer

A zone transfer can be defined as the process that occurs to copy the zone's resource records on the primary DNS server to secondary DNS servers. Zone transfer enables a secondary DNS server to continue handling queries if the primary DNS server fails. A secondary DNS server can also transfer its zone data to other secondary DNS servers that are beneath it in the DNS hierarchy. In this case, the secondary DNS server is regarded as the master DNS server to the other secondary servers.

The zone transfer methods are:

- *Full transfer*: When the user configures a secondary DNS server for a zone and starts the secondary DNS server, the secondary DNS server requests a full copy of the zone from the primary DNS server. A full transfer of all the zone information is performed. Full zone transfers tend to be resource intensive. This disadvantage of full transfers has led to the development of incremental zone transfers.
- *Incremental zone transfer*: With an incremental zone transfer, only those resource records that have since changed in a zone are transferred to the secondary DNS servers. During

zone transfer, the DNS database is on the primary.
DNS server and the secondary DNS server are compared to determine whether there are differences in the DNS data. If the primary and secondary DNS servers' data are the same, zone transfer does not take place. If the DNS data of the two servers are different, transfer of the delta resource records starts. This occurs when the serial number on the primary DNS server database is higher than that of secondary DNS server. For incremental zone transfer to occur, the primary DNS server has to record incremental changes to its DNS database. Incremental zone transfers require less bandwidth than full zone transfers.

- *Active Directory transfers*: These zone transfers occur when Active Directory-integrated zones are replicated to the domain controllers in a domain. Replication occurs through Active Directory replication.
- *DNS Notify* is a mechanism that enables a primary DNS server to inform secondary DNS servers when its database has been updated. DNS Notify informs the secondary DNS servers when they need to initiate a zone transfer so that the updates of the primary DNS server can be replicated to them. When a secondary DNS server receives the notification from the primary DNS server, it can start an incremental zone transfer or a full zone transfer to pull zone changes from the primary DNS servers.

Understanding DNS Resource Records (RRs)

The DNS database contains resource records (entries) that resolve name resolution queries sent to the DNS server. Each DNS server contains the resource records (RRs) it needs to respond to name resolution queries for the portion of the DNS namespace for which it is authoritative. There are different types of resource records.

A few of the commonly used resource records (RR) and their associated functions are described in the Table.

Resource Records Type	Name	Function
A	Host record	Contains the IP address of a specific host, and maps the FQDN addresses.
AAAA	IPv6 address record	Ties a FQDN to an IPv6 128-bit address.
AFSDB	Andrews files system	Associates a DNS domain name to a server subtype: an AFS v volume or an authenticated name server using DCE/NCA
ATMA	Asynchronous Transfer Mode address	Associates a DNS domain name to the ATM address of the atm_address field.
CNAME	Canonical Name / Alias name	Ties an alias to its associated domain name.
HINFO	Host info record	Indicates the CPU and OS type for a particular host.
ISDN	ISDN info record	Ties a FQDN to an associated ISDN telephone number
KEY	Public key resource record	Contains the public key for zones that can use DNS Security Extensions (DNSSEC).
MB	Mailbox name record	Maps the domain mail server name to the mail server.s host

		name
MG	Mail group record	Ties th domain mailing group to mailbox resource records
MINFO	Mailbox info record	Associates a mailbox for an individual that maintains it.
MR	Mailbox renamed record	Maps an older mailbox name to its new mailbox name.
MX	Mail exchange record	Provides routing for messages to mail servers and backup servers.
NS	Name server record	Provides a list of the authoritative servers for a domain. Also p the authoritative DNS server for delegated subdomains.
NXT	Next resource record	Indicates those resource record types that exist for a name. Sp the resource record in the zone.
OPT	Option resource record	A pseudo-resource record which provides extended DNS functionality.
PTR	Pointer resource record	Points to a different resource record, and is used for reverse lookups to point to A type resource records.
RT	Route through record	Provides routing information for hosts that do not have a WA address.
SIG	Signature resource record	Stores the digital signature for an RR set.
SOA	Start of Authority resource record	This resource record contains zone information for determining the name of the primary DNS server for the zone. zone property information, such as version information.
SRV	Service locator record	Used by Active directory to locate domain controllers, LDAP and global catalog servers.
TXT	Text record	Maps a DNS name to descriptive text.
X25	X.25 info record	Maps a DNS address to the public switched data network (PSI number.

While there are various resource records that contain different information, there are a few required fields that each particular resource record has to contain:

- *Owner* – the DNS domain that contains the resource record
- *TTL (Time to Live)* – indicates the time duration that DNS servers can cache resource record information prior to discarding the information. This is, however, an optional resource records field.
- *Class* – is another optional resource records field. Class types were used in earlier implementations of the DNS naming system and are no longer used these days.
- *Type* – indicates the type of information contained in the resource record.
- *Record Specific Data* – a variable length field that further defines the function of the resource. The format of the field is determined by Class and Type.

Delegation records and glue records can also be added to a zone. These records delegate a subdomain into a separate zone.

- *Delegation records*: These are Name Space (NS) resource records in a parent zone. The delegation record specifies the parent zone as being authoritative for the delegated zones.
- *Glue records*: These are A type resource records for the DNS server that has authority over delegated zone.

The more important resource records are discussed now. This includes the following:

- Start of Authority (SOA), Name Server (NS), Host (A), Alias (CNAME), Mail exchanger (MX), Pointer (PTR), Service location (SRV)

Start of Authority (SOA) Resource Record

This is the first record in the DNS database file. The SOA record includes information on the zone property information, such the primary DNS server for the zone and version information.

The fields located within the SOA record are listed below:

- *Source host* – the host for which the DNS database file is maintained
- *Contact e-mail* – e-mail address for the individual who is responsible for the database file.
- *Serial number* – the version number of the database.
- *Refresh time* – the time that a secondary DNS server waits while determining whether database updates have been made that have to be replicated via zone transfer.
- *Retry time* – the time for which a secondary DNS server waits before attempting a failed zone transfer again.
- *Expiration time* – the time for which a secondary DNS server will continue to attempt to download zone information. Old zone information is discarded when this limit is reached.
- *Time to live* – the time that the particular DNS server can cache resource records from the DNS database file.

Name Server (NS) Resource Record

The Name Server (NS) resource record provides a list of the authoritative DNS servers for a domain as well authoritative DNS server for any delegated subdomains. Each zone must have one (or more) NS resource records at the zone root. The NS resource record indicates the primary and secondary DNS servers for the zone defined in the SOA resource record. This in turn enables other DNS servers to look up names in the domain.

Host (A) Resource Record

The host (A) resource record contains the IP address of a specific host and maps the FQDN to this 32-bit IPv4 addresses. Host (A) resource records basically associates the domain names of computers (FQDNs) or hosts names to their associated IP addresses. Because a host (A) resource record statically associates a host name to a specific IP address, users can manually add these records to zones if they have machines that have statically assigned IP addresses.

The methods used to add host (A) resource records to zones are:

- Manually add these records using the DNS management console.
- Use the Dnscmd tool at the command line to add host (A) resource records.
- TCP/IP client computers running Windows 2000, Windows XP, or Windows Server 2003 use the DHCP Client service to both register their names and update their host (A) resource records.

Alias (CNAME) Resource Record

Alias (CNAME) resource records tie an alias name to its associated domain name. Alias (CNAME) resource records are referred to as *canonical names*. By using canonical names, users can hide network information from the clients connected to their network. Alias (CNAME) resource records should be used when users have to rename a host that is defined in a host (A) resource record in the identical zone.

Mail Exchanger (MX) Resource Record

The mail exchanger (MX) resource record provides routing for messages to mail servers and backup servers. The mail MX resource record provides information on which mail servers processes e-mail for the particular domain name. E-mail applications therefore mostly utilize MX resource records.

A mail exchanger (MX) resource record has the following parameters:

- Priority
- Mail server

The mail exchanger (MX) resource record enables the DNS server to work with e-mail addresses where no specific mail server is defined. A DNS domain can have multiple MX records. MX resource records can therefore also be used to provide failover to different mail servers when the primary server specified is unavailable. In this case, a server preference value is added to indicate the priority of a server in the list. Lower server preference values specify higher preference.

Pointer (PTR) Resource Record

The pointer (PTR) resource record points to a different resource record and is used for reverse lookups to point to A resource records. Reverse lookups resolve IP addresses to host names or FQDNs.

Add PTR resource records to zones through the following methods:

- Manually add these records with the DNS management console.
- Use the Dnscmd tool at the command line to add PTR resource records.

Service (SRV) Resource Records

Service (SRV) resource records are typically used by Active directory to locate domain controllers, LDAP servers, and global catalog servers. The SRV records define the location of specific services in a domain. They associate the location of a service such as a domain controller or global catalog server with details on how the particular service can be contacted.

The fields of the service (SRV) resource record are explained below:

- Service name
- The protocol used
- The domain name associated with the SRV records
- The port number for the particular service
- The Time to Live value
- The class
- The priority and weight
- The target specifying the FQDN of the particular host supporting the service

The Zone Database Files

If the user is not using Active Directory-integrated zones, the specific zone database files that are used for zone data are:

- *Domain Name file*: When new A type resource records are added to the domain, they are stored in this file. When a zone is created, the Domain Name file contains the following:
    - An SOA resource record for the domain
    - An NS resource record that indicates the name of the DNS server that was created.
- *Reverse Lookup file*: This database file contains information on a reverse lookup zone.
- *Cache file*: This file contains a listing of the names and addresses of root name servers that are needed for resolving names that are external to the authoritative domains.
- *Boot file*: This file controls the DNS server's startup behavior. The boot file supports the commands listed below:
    - Directory command – this command defines the location of the other files specified in the Boot file.
    - Primary command – defines the domain for which this particular DNS server has authority.
    - Secondary – specifies a domain as being a secondary domain.
    - Cache command – this command defines the list of root hints used for contacting DNS servers for the root domain.

Planning DNS Zone Implementations

When users divide up the DNS namespace, DNS zones are created. Breaking up the namespace into zones enables DNS to more efficiently manage available bandwidth usage, which in turn improves DNS performance.

When *determining how to break up the DNS zones*, a few considerations to take include:

- DNS traffic patterns: use the System Monitor tool to examine DNS performance counters and to obtain DNS server statistics.
- Network link speed: The types of network links that exist between DNS servers should be determined when users plan the zones for their environment.
- Whether full DNS servers or caching-only DNS servers are being used also affects how users break up DNS zones.

The main zone types used in Windows Server 2003 DNS environments are primary zones and Active Directory-integrated zones. The question on whether to implement primary zones or Active Directory-integrated zones would be determined by the environment's DNS design requirements.

Both primary zones and secondary zones are standard DNS zones that use zone files. The main difference between primary zones and secondary zones is that primary zones can be updated. Secondary zones contain read-only copies of zone data. A secondary DNS zone can only be updated through DNS zone transfer. Secondary DNS zones are usually implemented to provide fault tolerance for the DNS server environment.

An Active Directory-integrated zone can be defined as an improved version of a primary DNS zone because it can use multi-master replication and the security features of Active Directory. The zone data of Active Directory-integrated zones are stored in Active Directory. Active Directory-integrated zones are authoritative primary zones.

A few advantages that Active Directory-integrated zone implementations have over standard primary zone implementations are:

- Active Directory replication is faster, which means that the time needed to transfer zone data between zones is far less.
- The Active Directory replication topology is used for Active Directory replication and for Active Directory-integrated zone replication. There is no longer a need for DNS replication when DNS and Active Directory are integrated.
- Active Directory-integrated zones can enjoy the security features of Active Directory.
- The need to manage Active Directory domains and DNS namespaces as separate entities is eliminated. This in turn reduces administrative overhead.
- When DNS and Active Directory are integrated, the Active Directory-integrated zones are replicated and stored on any new domain controllers automatically. Synchronization takes place automatically when new domain controllers are deployed.

The mechanism that DNS utilizes to forward a query that one DNS server cannot resolve to another DNS server is called *DNS forwarding*. *DNS forwarders* are the DNS servers used to

forward DNS queries for different DNS namespace to those DNS servers who can answer the query. A DNS server is configured as a DNS forwarder when users configure the other DNS servers to direct any unresolved queries to a specific DNS server. Creating DNS forwarders can improve name resolution efficiency.

Windows Server 2003 DNS introduces a new feature called *conditional forwarding*. With conditional forwarding, users create *conditional forwarders* within their environment that will forward DNS queries based on the specific domain names being requested in the query. This differs from DNS forwarders where the standard DNS resolution path to the root was used to resolve the query. A conditional forwarder can only forward queries for domains that are defined in the particular conditional forwarders list. The query is passed to the default DNS forwarder if there are no entries in the forwarders list for the specific domain queried.

When conditional forwarders are configured, the process to resolve domain names is illustrated below:

1. A client sends a query to the DNS server for name resolution.
2. The DNS server checks its DNS database file to determine whether it can resolve the query with its zone data.
3. The DNS server also checks its DNS server cache to resolve the request.
4. If the DNS server is not configured to use forwarding, the server uses recursion to attempt to resolve the query.
5. If the DNS server is configured to forward the query for a specific domain name to a DNS forwarder, the DNS server then forwards the query to the IP address of its configured DNS forwarder.

A few *considerations for configuring forwarders for the DNS environment* are:

- Only implement the DNS forwarders that are necessary for the environment. Refrain from creating loads of forwarders for the internal DNS servers.
- Avoid chaining your DNS servers together in a forwarding configuration.
- To avoid the DNS forwarder turning into a bottleneck, do not configure one external DNS forwarder for all the internal DNS servers.

How to Create a New Zone

1. Click Start, Administrative Tools, and DNS to open the DNS console.
2. Expand the Forward Lookup Zones folder.
3. Select the Forward Lookup Zones folder.
4. From the Action menu, select New Zone.
5. The New Zone Wizard initiates.
6. On the initial page of the Wizard, click Next.
7. On the Zone Type page, ensure that the Primary Zone Creates A Copy Of A Zone That Can Be Updated Directly On This Server option is selected. This option is selected by default.

8. Uncheck the Store The Zone In Active Directory (Available Only If DNS Server Is A Domain Controller) checkbox.
   Click Next.
9. On the Zone Name page, enter the correct name for the zone in the Zone Name textbox. Click Next.
10. On the Zone File page, ensure that the default option, Create A New File With This File Name, is selected. Click Next.
11. On the Dynamic Update page, ensure that the Do Not Allow Dynamic Updates. Dynamic Updates Of Resource Records Are Not Accepted By This Zone. You Must Update These Records Manually option is selected. Click Next.
12. The Completing The New Zone Wizard page is displayed next.
13. Click Finish to create the new zone.

How to Create Subdomains

1. Click Start, Administrative Tools, and DNS to open the DNS console.
2. In the console tree, select the appropriate zone.
3. From the Action menu, select New Domain.
4. The DNS Domain dialog box opens.
5. Enter the name for new subdomain.
6. Click OK to create the new subdomain.

How to create a reverse lookup zone

1. Click Start, Administrative Tools, and DNS to open the DNS console.
2. Select the appropriate DNS server in the console tree.
3. Right-click the DNS server then select New Zone from the shortcut menu.
4. The New Zone Wizard starts.
5. Click Next on the first page of the New Zone Wizard.
6. On the Zone Type page, ensure that the Primary Zone option is selected. Click Next.
7. On the following page, select the Reverse lookup zone option. Click Next.
8. Enter the IP network in the Network ID box for the domain name that the new reverse lookup zone is being created for. Click Next.
9. Accept the default zone file name. Click Next.
10. On the Dynamic Update page, select the Allow both nonsecure and secure dynamic updates option, then click Next.
11. The Completing The New Zone Wizard page is displayed next.
12. Click Finish to create the new reverse lookup zone.

How to Create a Stub Zone

1. Click Start, Administrative Tools, and then click DNS to open the DNS console.
2. Expand the Forward Lookup Zones folder.
3. Select the Forward Lookup Zones folder.
4. From the Action menu, select New Zone.

5. The New Zone Wizard initiates.
6. On the initial page of the Wizard, click Next.
7. On the Zone Type page, select the Stub Zone option.
8. Uncheck the Store The Zone In Active Directory (Available Only If DNS Server Is A Domain Controller) checkbox. Click Next.
9. On the Zone Name page, enter the name for the new stub zone in the Zone Name textbox then click Next.
10. Accept the default setting on the Zone file page. Click Next.
11. On the Master DNS Servers page, enter the IP address of the master server in the Address text box. Click Next.
12. On the Completing The New Zone Wizard page, click Finish.

How to Add Resource Records to Zones

1. Click Start, Administrative Tools, and DNS to open the DNS console.
2. In the console tree, select the zone to add resource records to.
3. From the Action menu, select the resource record type to be added to the zone. The options are:
    - New Host (A)
    - New Alias (CNAME)
    - New Mail Exchanger (MX)
    - Other New Records
4. Select the New Host (A) option.
5. The New Host dialog box opens.
6. In the Name (Use Parent Domain Name If Blank) textbox, enter the name of the new host.
7. When the user specifies the name of the new host, the resulting FQDN is displayed in the Fully qualified domain name (FQDN) textbox.
8. In the IP Address box, enter the address for the new host.
9. To create an associated pointer (PTR) record, enable the checkbox.
10. Click the Add Host button.
11. The new host (A) resource record is added to the particular zone.
12. A message box is displayed, verifying that the new host (A) resource record was successfully created for the zone.
13. Click OK.
14. Click Done to close the New Host dialog box./li>

How to Create a Zone Delegation

1. Click Start, Administrative Tools, and select DNS to open the DNS console.
2. Right-click the subdomain in the console tree, then select New Delegation from the shortcut menu.
3. The New Delegation Wizard initiates.
4. Click Next on the first page of the New Delegation Wizard.

5. When the Delegated Domain Name page opens, provide a delegated domain name then click Next.
6. On the Name Servers page, click the Add button to provide the DNS servers' names and IP addresses that should host the delegation.
7. On the Name Servers page, click Next.
8. Click Finish.

How to Enable Dynamic Updates for a Zone

1. Click Start, Administrative Tools, and the select DNS to open the DNS console.
2. Right-click the zone to work with in the console tree, then select Properties from the shortcut menu.
3. When the Zone Properties dialog box opens, on the General tab, select Yes in the Allow Dynamic Updates list box.
4. Click OK.

How to Configure a Zone to Use WINS for Name Resolution

Users can configure their forward lookup zone to use WINS for name resolution in instances where the queried name is not found in the DNS namespace.

1. Click Start, Administrative Tools, and DNS to open the DNS console.
2. In the console tree, expand the DNS server node then expand the Forward Lookup Zones folder.
3. Locate and right-click the zone to be configured, then select Properties from the shortcut menu.
4. When the Zone Properties dialog box opens, click the WINS tab.
5. Enable the Use WINS Forward Lookup checkbox.
6. Type the WINS server IP address. Click Add, then OK.
7. On the General tab, select Yes in the Allow Dynamic Updates list box.
8. Click OK.

Related Articles on DNS

- What is DNS?
- How do I flush DNS?
- How do I find my DNS servers?
- What are public DNS servers?
- How do I perform a DNS lookup?
- What is reverse DNS?
- What is a dynamic DNS?
- What are DNS root servers?
- Understanding DNS