Replication Between Sites

The main characteristics about replication within sites are as follows:

- The network connections within a site are both reliable, cheap, and have sufficient available bandwidth.
- Replication traffic within a site is not compressed, because a site assumes fast, highly reliable network connections. Not compressing replication traffic helps reduce the processing load on the domain controllers. However, uncompressed traffic may increase the network bandwidth.
- A change notification process initiates replication within a site.

The main characteristics about replication between sites are as follows:

- The network links between sites have limited available bandwidth, may have a higher cost, and may not be reliable.
- Replication traffic between sites can be designed to optimize bandwidth by compressing all replication traffic. Replication traffic is compressed to 10 to 15 percent of its original size before it is transmitted. Although compression optimizes network bandwidth, it imposes an additional processing load on domain controllers when it compresses and decompresses replication data.
- Replication between sites occurs automatically after you have defined configurable values, such as a schedule or a replication interval. You can schedule replication for inexpensive or off-peak hours. By default, changes are replicated between sites according to a schedule that you define, and not according to when changes occur. The schedule determines when replication can occur. The interval specifies how often domain controllers check for changes during the time that replication can occur.

Change Notifications Between AD DS Sites

By design, changes in AD DS replicate between domain controllers in different sites according to a defined replication schedule, and not according to when changes occur, such as with intrasite replication. Because of this, the replication latency in the forest can equal the sum of the greatest replication latencies along the longest replication path of any directory partition. In some scenarios, this can be inefficient.

To avoid latency in replication, you can configure change notifications on connections between sites. By modifying the site link object, you can enable change notification between sites for all connections that occur over that link. Since the replication partner across the site is notified of changes, the intersite replication interval is effectively ignored. The originating domain controller notifies the domain controller in the other site that it has a change, just as it does within a single site.

For changes such as account locks or similar security-related changes, immediate replication is crucial. In these situations, urgent replication is used. Urgent replication bypasses the notification delay and processes the change notifications immediately. This only affects change notifications. If you do not have change notifications enabled between sites, replication still honors the replication interval on the site link.

Note: When the user's password is changed, immediate replication is initiated to the primary domain controller (PDC) emulator operations master. This differs from urgent replication because it occurs immediately without regard to the inter-site replication interval.

What Is the Intersite Topology Generator?

When you configure multiple sites, the KCC on one domain controller in each site is designated as the site's intersite topology generator (ISTG). There is only one ISTG per site, regardless of how many domains or other directory partitions the site has. ISTG is responsible for calculating the site's ideal replication topology.

When you add a new site to the forest, each site's ISTG determines which directory partitions are present in the new site. The ISTG then calculates how many new connection objects are necessary to replicate the new site's required information. In some networks, you might want to specify that only certain domain controllers are responsible for intersite replication. You can do this by specifying bridgehead servers. The bridgehead servers are responsible for all replication into, and out of, the site. ISTG creates the required connection agreement in its directory, and this information is then replicated to the bridgehead server. The bridgehead server then creates a replication connection with the bridgehead server in the remote site, and replication begins. If a replication partner becomes unavailable, the ITSG selects another domain controller automatically, if possible. If bridgehead servers have been assigned manually, and if they become unavailable, ISTG will not automatically select other servers.

The ISTG selects bridgehead servers automatically, and creates the intersite replication topology to ensure that changes replicate effectively between bridgeheads that share a site link. Bridgeheads are selected per partition, so it is possible that one domain controller in a site might be the bridgehead server for the schema, while another is for the configuration. However, you usually will find that one domain controller is the bridgehead server for all partitions in a site, unless there are domain controllers from other domains or application directory partitions. In this scenario, bridgeheads will be chosen for those partitions.

Overview of SRV Resource Records for Domain Controllers

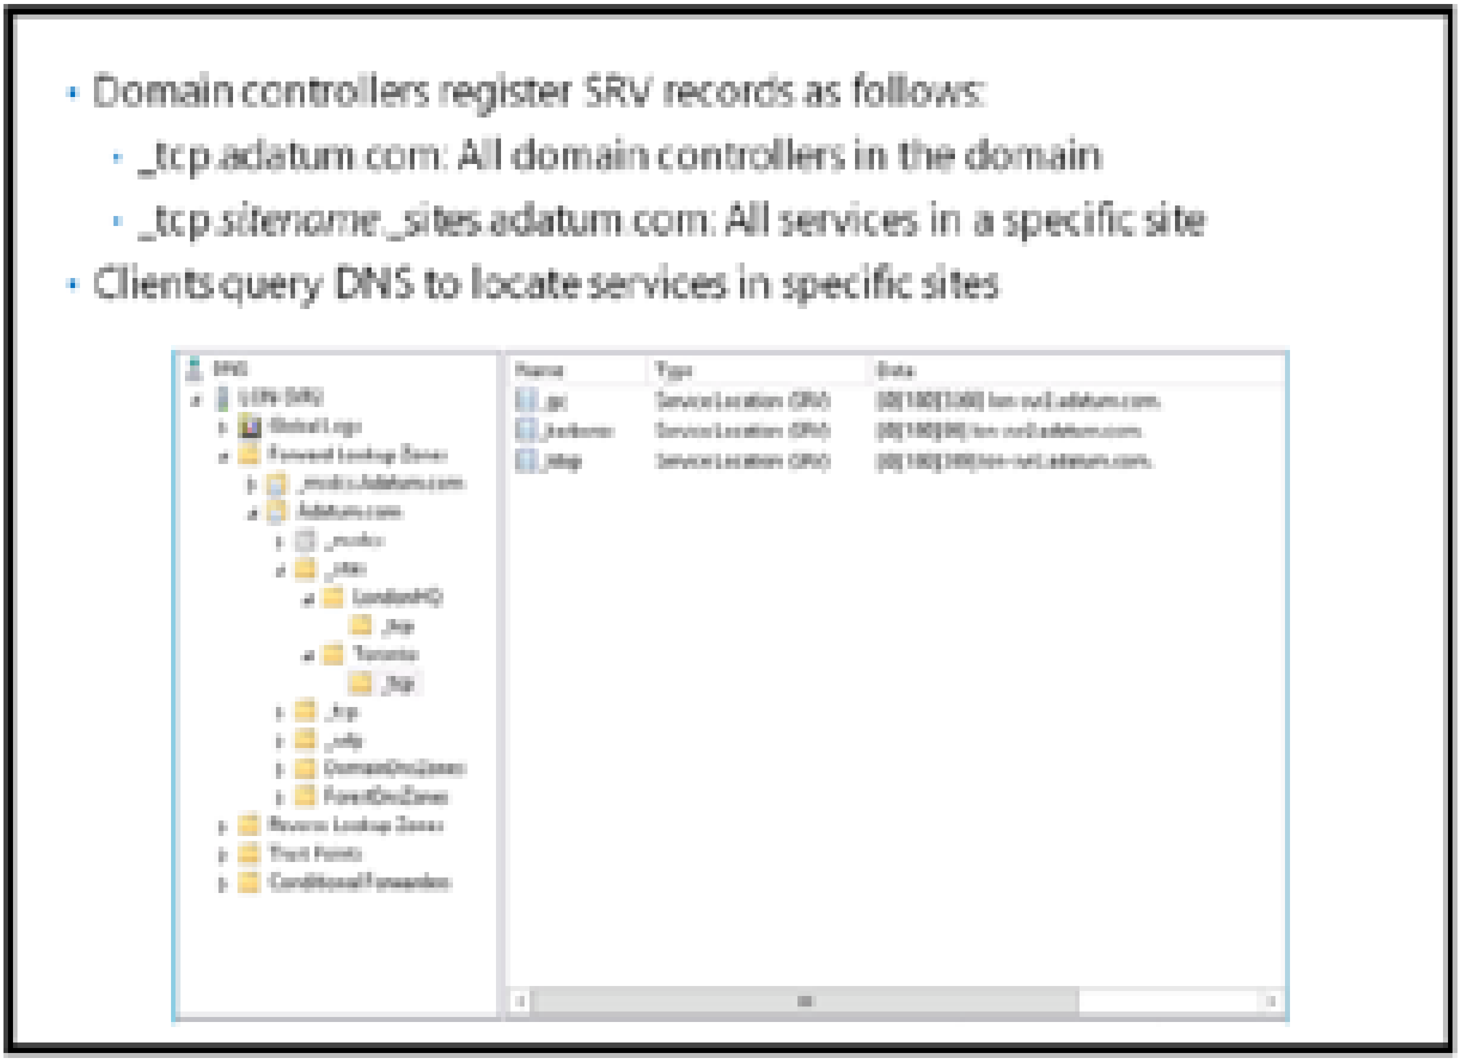

- Domain controllers register SRV records as follows:
  - _tcp.adatum.com: All domain controllers in the domain
  - _tcp.*sitename*._sites.adatum.com: All services in a specific site
- Clients query DNS to locate services in specific sites

Name	Type	Data
_gc	Service Location (SRV)	[0][100][3268] lon-svr1.adatum.com.
_kerberos	Service Location (SRV)	[0][100][88] lon-svr1.adatum.com.
_ldap	Service Location (SRV)	[0][100][389] lon-svr1.adatum.com.

When you add a domain controller to a domain, the domain controller advertises its services by creating service (SRV) resource records (also known as *locator records*) in DNS. Unlike host (A) resource records, which map host names to IP addresses, SRV records map services to host names. For example, to publish its ability to provide authentication and directory access, a domain controller registers Kerberos v5 protocol and LDAP SRV records. These SRV records are added to several folders within the forest's DNS zones.

Within the domain zone, a folder called *name*_tcp contains the SRV records for all domain controllers in the domain. Additionally, within the domain zone is a folder called *name*_sites, which contains subfolders for each site configured in the domain. Each site-specific folder contains SRV records that represent services available in the site. For example, if a domain controller is located in a site, a SRV record will be located at the path _sites\sitename_tcp, where *sitename* is the name of the site.

A typical SRV record contains the following information:

- The service name and port. This portion of the SRV record indicates a service with a fixed port. It does not have to be a well-known port. SRV records in Windows Server 2012 include LDAP (port 389), Kerberos (port 88), Kerberos password protocol (KPASSWD, port 464), and global catalog services (port 3268).
- Protocol. The Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) is indicated as a transport protocol for the service. The same service can use both protocols in separate SRV records. Kerberos records, for example, are registered for both TCP and UDP. Microsoft clients use only TCP, but UNIX clients can use both UDP and TCP.
- Host name. The host name corresponds to the host A record for the server hosting the service. When a client queries for a service, the DNS server returns the SRV record and associated host A records, so the client does not need to submit a separate query to resolve the IP address of a service.

The service name in an SRV record follows the standard DNS hierarchy with components separated by dots. For example, a domain controller’s Kerberos service is registered as: kerberos._tcp.*sitename*._sites.*domainname*, where:

- *domainName* is the domain or zone, for example contoso.com.
- _sites is all sites registered with DNS.
- *sitename* is the site of the domain controller registering the service.
- _tcp is any TCP-based services in the site.
- kerberos is a Kerberos Key Distribution Center (KDC)that uses TCP as its

transport protocol.

How Client Computers Locate Domain Controllers Within Sites

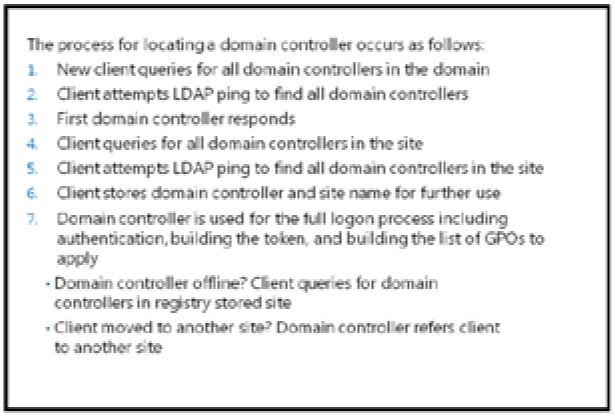

The process for locating a domain controller occurs as follows:

1. New client queries for all domain controllers in the domain
2. Client attempts LDAP ping to find all domain controllers
3. First domain controller responds
4. Client queries for all domain controllers in the site
5. Client attempts LDAP ping to find all domain controllers in the site
6. Client stores domain controller and site name for further use
7. Domain controller is used for the full logon process including authentication, building the token, and building the list of GPOs to apply

- Domain controller offline? Client queries for domain controllers in registry stored site
- Client moved to another site? Domain controller refers client to another site

When you join a Windows operating system client to a domain and then restart it, the client completes a domain controller location and registration process. The goal of this registration process is to locate the domain controller with the most efficient and closest location to the client's location based on IP subnet information.

The process for locating a domain controller is as follows:

1. The new client queries for all domain controllers in the domain. As the new domain client restarts, it receives an IP address from a DHCP server, and is ready to authenticate to the domain. However, the client does not know where to find a domain controller. Therefore, the client queries for a domain controller by querying the _tcp folder, which contains the SRV records for all domain controllers in the domain.

2. The client attempts an LDAP ping to all domain controllers in a sequence. DNS returns a list of all matching domain controllers, and the client attempts to contact all of them on its first startup.

3. The first domain controller responds. The first domain controller that responds to the client examines the client's IP address, cross-references that address with subnet objects, and informs the client of the site to which the client

belongs. The client stores the site name in its registry, and then queries for domain controllers in the site-specific _tcp folder.

4. The client queries for all domain controllers in the site. DNS returns a list of all domain controllers in the site.

5. The client attempts an LDAP ping sequentially to all domain controllers in the site. The domain controller that responds first authenticates the client.

6. The client forms an affinity. The client forms an affinity with the domain controller that responded first, and then attempts to authenticate with the same domain controller in the future. If the domain controller is unavailable, the client queries the site's _tcp folder again, and again attempts to bind with the first domain controller that responds in the site.

If the client moves to another site, such as the case for a mobile computer, the client attempts to authenticate to its preferred domain controller. The domain controller notices that the client's IP address is associated with a different site, and then refers the client to the new site. The client then queries DNS for domain controllers in the local site

Automatic Site Coverage

As mentioned previously, you can configure sites to direct users to local copies of replicated resources, such as shared folders replicated within a DFS namespace. There may be scenarios in which you only require service localization with no need for a domain controller located within the site. In this case, a nearby domain controller will register its SRV records in the site by using a process called site coverage.

A site without a domain controller generally is covered by a domain controller in a site with the lowest site-link cost to the site that requires coverage. You also can configure site coverage and SRV record priority manually if you want to control authentication in sites without domain controllers.

Additional Reading: For more information about how site coverage is evaluated, see http://go.microsoft.com/fwlink/?LinkId=168550.