Introduction to configuring advanced replication in the IBM Tivoli Directory Server for z/OS

z/OS Operating System V1R11

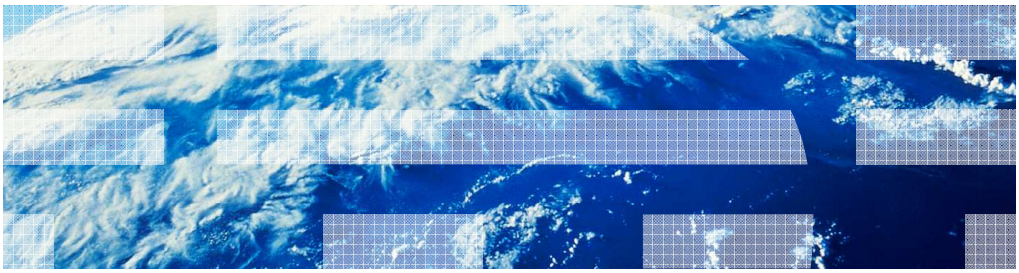

This presentation briefly describes the advanced replication features and how to configure advanced replication in the IBM Tivoli Directory Server for z/OS. It is applicable to customers running at z/OS 1.11 and above.

Agenda

- Types of advanced replication servers and roles
- Overview of advanced replication features
- Types of replication topologies
- Replication topology entries
- Configuring advanced replication
- Setting up a master-replica topology

This presentation presents an overview of the advanced replication features, the types of advanced replication servers, and the different types of advanced replication topologies. It also discusses the replication topology entries and how to configure a master-replica topology.

Types of advanced replication servers

- Consumer – A server that receives replication changes from another (supplier) server
- Forwarder – A read-only consumer server that replicates all changes sent to it
- Gateway – A server that forwards all replication traffic from the local replication site where it resides to other gateway servers in the replicating network
- Master – A supplier server that is writable for a given subtree
- Peer – A supplier server where there are multiple masters for a given subtree
- Supplier – A server that sends replication changes to another (consumer) server

Advanced replication has introduced some new server roles. The terminology consumer and supplier indicates whether a server is a receiver or sender of advanced replication updates. When a server is a supplier, it sends replication changes to another server called a consumer. In advanced replication, each subtree is replicated independently and therefore each server can have multiple roles. A server can be configured to be a consumer server for subtree o=xyz while being a supplier server for subtree o=abc.

A peer is a supplier server where there are multiple masters or writable servers for a given subtree.

The forwarder and gateway server roles are new in advanced replication. A forwarder or cascading server is a read only server that replicates all updates that are sent to it from a supplier server. A gateway server forwards all replication traffic from a local replication site to other gateway servers in other replication sites.

Advanced replication features

- Subtree-based replication – Allows only a portion of the DIT to be replicated instead of replicating an entire backend
- Schema replication – Replication of the schema entry
- Filtered/partial replication – Enables attributes and updates from being replicated to a consumer server
- Scheduled replication – Can schedule replication to occur during off-peak hours
- Conflict resolution – Provides automatic correction of entries in case updates arrive out of order on consumer servers

Unlike basic replication which requires an entire backend to be replicated, advanced replication does its replication based upon subtrees. The subtrees that are allowed to be replicated are configurable in advanced replication.

Schema replication is now supported with advanced replication. This allows the schema entry to be kept automatically in sync between supplier and consumer servers.

Advanced replication supports filtered or partial replication which includes or excludes entries or attribute types from being replicated from a supplier to a consumer server.

Replication can now be scheduled to occur during off-peak hours by configuring the days of the week and the times during the day it is allowed to occur.

Advanced replication provides automatic conflict resolution when updates arrive out of order in a consumer server. This support helps to prevent the supplier and consumer servers from getting out of sync when there are multiple supplier servers connected to the same consumer server.

Advanced replication features (continued)

- Extended operations allow the LDAP administrator to maintain and administer replication configurations and recover from replication related errors
- Operational attributes on replication topology entries provide current status

Advanced replication supports a number of advanced replication extended operations to assist the LDAP administrator in monitoring the advanced replication environment. Among the features these extended operations can perform are: suspend or resume replication, show replication errors, and skip or remove replication errors from the replication queue.

Along with using the extended operations, the LDAP administrator can now use LDAP search commands to obtain operational attribute values. These operational attribute values give the LDAP administrator current status information on the state of the advanced replication environment.

Master-replica topology

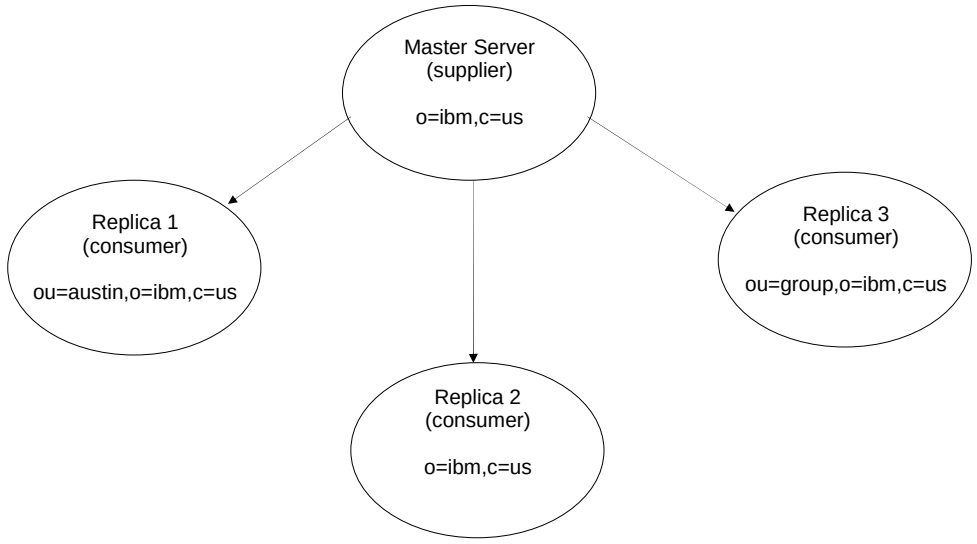

This slide shows an example of a master-replica topology in advanced replication. The master server accepts both client read and update operations. However, a replica server only accepts client read operations. If a client sends update requests to a replica server, the requests are referred to the master server which performs the updates and then replicates them to the replica server.

In the example shown on this slide, the master server has been configured to replicate updates under the o=ibm,c=us subtree to one of three different replica servers. If an entry under the o=ibm,c=us subtree is updated on the master server, the updates are only replicated to replica server 2. If an entry under the ou=austin,o=ibm,c=us subtree is updated on the master server, the updates are replicated to replica servers 1 and 2. If an entry under the ou=group,o=ibm,c=us subtree is updated on the master server, the updates are replicated to replica servers 2 and 3.

Peer-to-peer topology

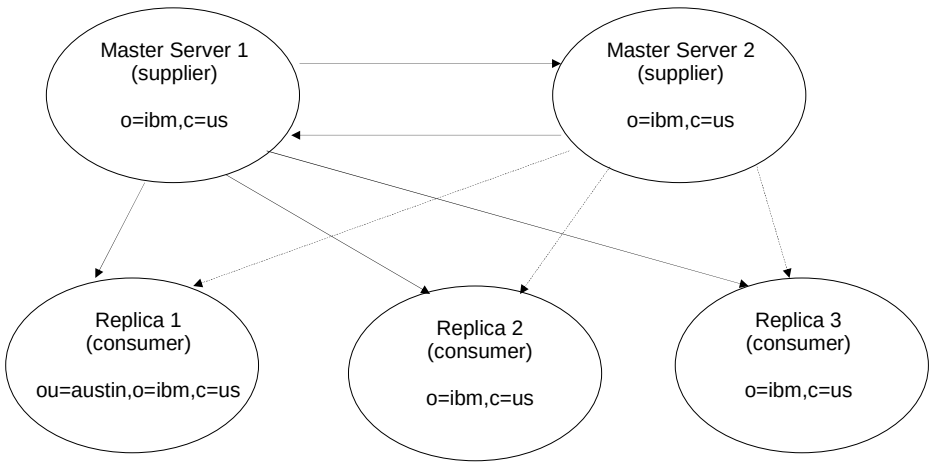

This slide shows an example of a peer-to-peer topology in advanced replication. In a peer-to-peer replication topology, there are several servers acting as masters for directory information, with each master responsible for updating other master servers and replica servers. A peer-to-peer replication topology can improve performance, availability, and reliability. Performance is improved by providing a local server to handle updates in a widely distributed network. Availability and reliability are improved by providing a backup master server ready to take over immediately if the primary master fails. Peer master servers replicate all client updates to the replicas and to the other peer masters, but do not replicate updates received from other master servers.

In the example shown on this slide, there are two master servers which are configured to replicate updates under the o=ibm,c=us subtree. If master server 1 receives an update client request under the o=ibm,c=us subtree, the updates are replicated to master server 2 and replica servers 2 and 3. Note, master server 2 will not replicate updates to replica servers 2 and 3 because the update was received from master server 1.

Similarly, if master server 2 receives an update request under the o=ibm,c=us subtree, the updates are replicated to master server 1 and replica servers 2 and 3. Note, master server 1 will not replicate updates to replica servers 2 and 3 because the update was received from master server 2.

Forwarding (Cascading) topology

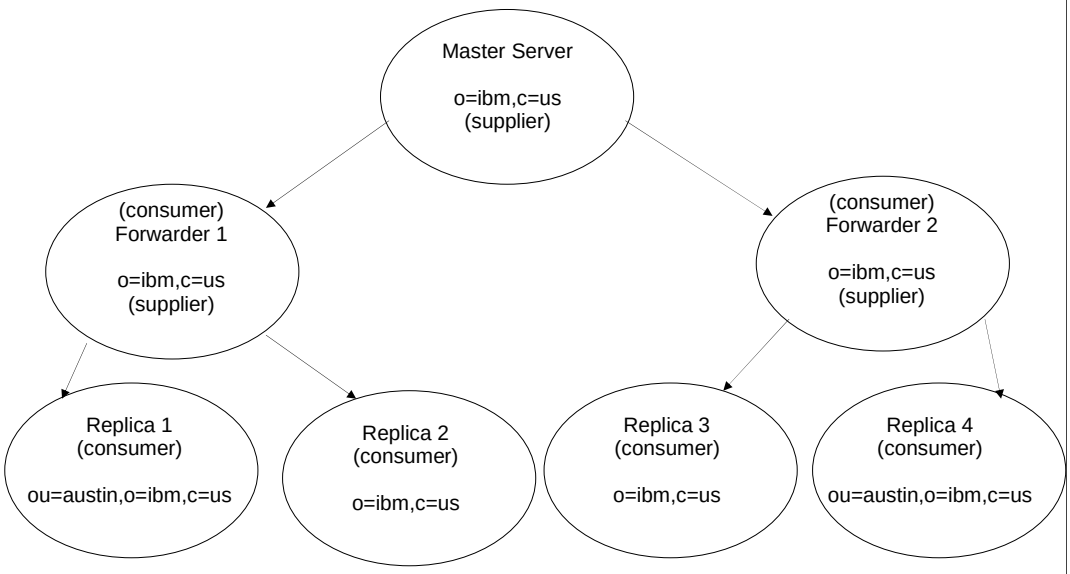

A forwarding or cascading replication is a replication topology that has multiple tiers of servers. A master server replicates to a set of read-only or forwarding servers that in turn replicate to other servers. A forwarding or cascading topology off-loads replication work from the master server. If a client sends update requests to a forwarding server, the requests are referred to the master server which performs the updates and then replicates them to the forwarding or cascading server.

In the example shown on this slide, the master server is a supplier server to the two forwarding servers. The forwarding servers serve two roles. They are consumer servers of the master server and supplier servers to the replica servers associated with them. The replica servers are consumer servers of their respective forwarding servers. If the master server receives an update client request under the o=ibm,c=us subtree, the updates are replicated to the forwarder 1 and 2 servers. Forwarder 1 and 2 servers then replicate the updates to replica servers 2 and 3.

Gateway topology

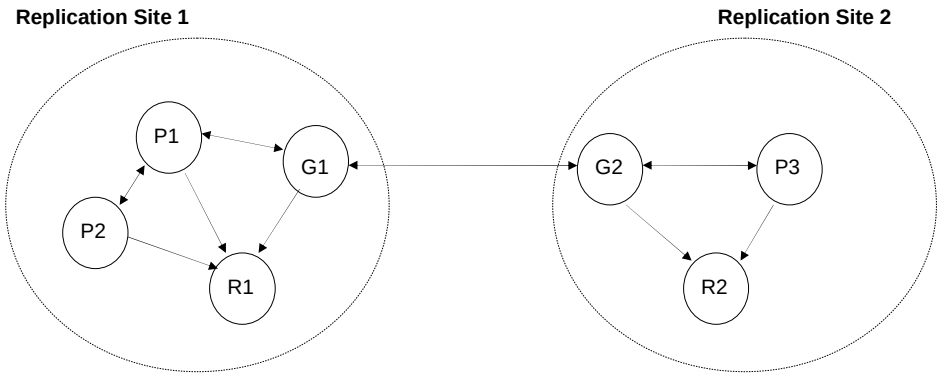

P = Peer server
G = Gateway server
R = Read-only replica server

A gateway replication topology is a more complex adaptation of peer-to-peer replication that extends replication capabilities across networks. The most notable difference is that a gateway server replicates changes received from other peer servers through the gateway. A gateway server must be a master server which accepts client update and read requests. It acts as a peer server within its own replication site. That is, it can receive and replicate client updates and receive updates from the other peer-master servers within the replication site. It does not replicate the updates received from the other peer-masters to any servers within its own site.

Within the gateway network, the gateway server acts as a two-way forwarding server. In one instance, the peer servers in its replication site act as the supplier servers to its own gateway server and the gateway servers at other replication sites are its consumer servers. In the other instance the situation is reversed. The other gateway servers are suppliers to the local gateway server and the servers within the local gateway's replication site are consumer servers.

Gateway replication uses gateway servers to collect and distribute replication information effectively across a replicating network. The primary benefit of gateway replication is the reduction of network traffic.

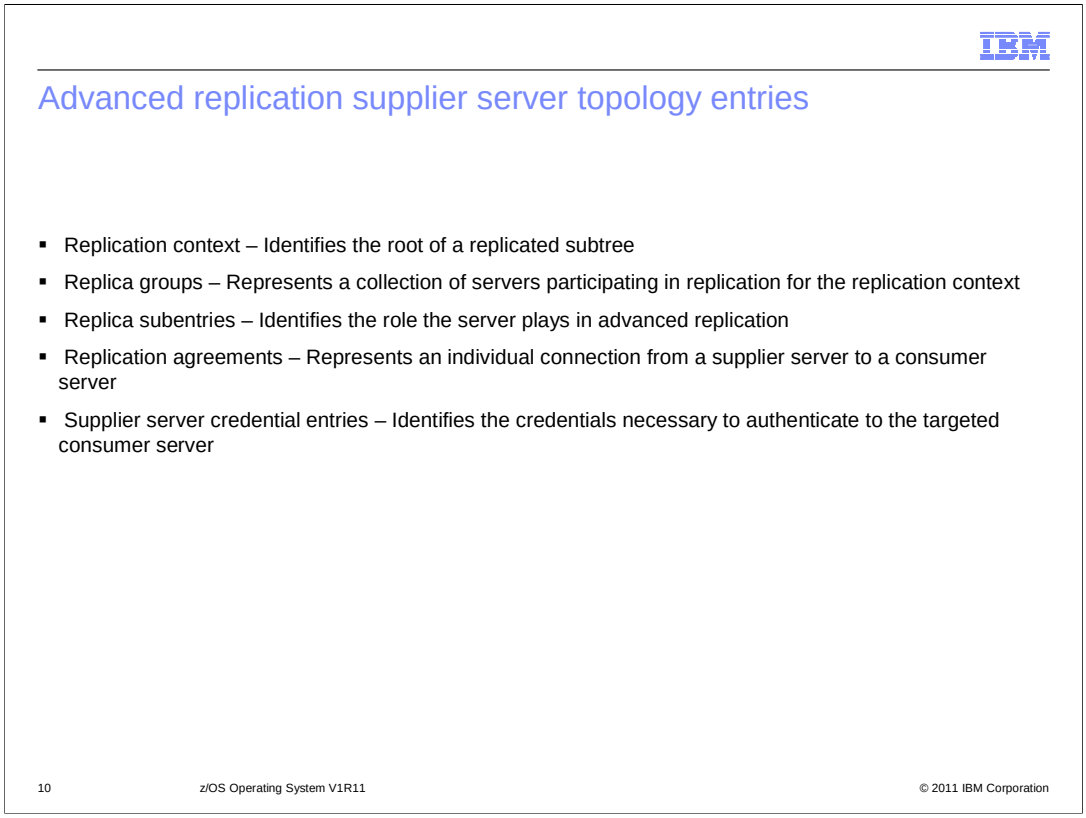

This slide briefly describes the entries that are needed on the supplier server when configuring advanced replication.

A replication context entry identifies the root of a replicated subtree for advanced replication. Created directly under a replication context entry is a replica group entry which represents a collection of servers participating in replication for the context. Multiple replica group entries are allowed to be created under a replication context. A replica subentry is created directly under a replica group and identifies the role the server plays in advanced replication (for example master or read-only replica).

A replication agreement entry is created directly under a replica subentry to define replication from the server represented by the subentry to another server. This entry represents an individual connection from a supplier server to a consumer server. A replica subentry may have any number of replication agreement entries defined under it to specify each supplier agreement this server has under this replication topology.

Because the replication agreement entry can be replicated, a distinguished name (DN) to credentials object is used in the **ibm-replicaCredentialsDN** attribute value of the replication agreement entry. This allows the supplier server credentials entry to be stored in an area of the DIT that is not replicated. Replicating the supplier server credentials entries (where 'clear text' passwords must be obtainable) represents a potential security exposure.

Advanced replication supplier server topology entries (continued)

- Replication filter – Indicates the entries and attribute types that are allowed to be replicated to a consumer server
- Replication schedule – Specifies the time of day and day of the week when replication is allowed to occur

The replication filter and replication schedule entries are optional supplier server entries. A replication filter entry allows the LDAP administrator to include or exclude certain entries or attribute types from being replicated to the consumer server for each individual replication agreement.

A replication schedule entry allows the LDAP administrator to schedule advanced replication to occur at optimal times when network traffic is minimal for each individual replication agreement. The replication schedule can be configured to specify the times during the day and days of the week when replication must occur to the consumer server.

Advanced replication consumer server topology entry

- Consumer server credential entries – Identifies the credentials that the supplier server is using to authenticate to the consumer server

The only replication related entry needed on the consumer server is the consumer server credentials entry, which must reside under the **cn=configuration** suffix in the CDBM backend. The consumer server credentials entry specifies the bind DN and password of the supplier server performing replication.

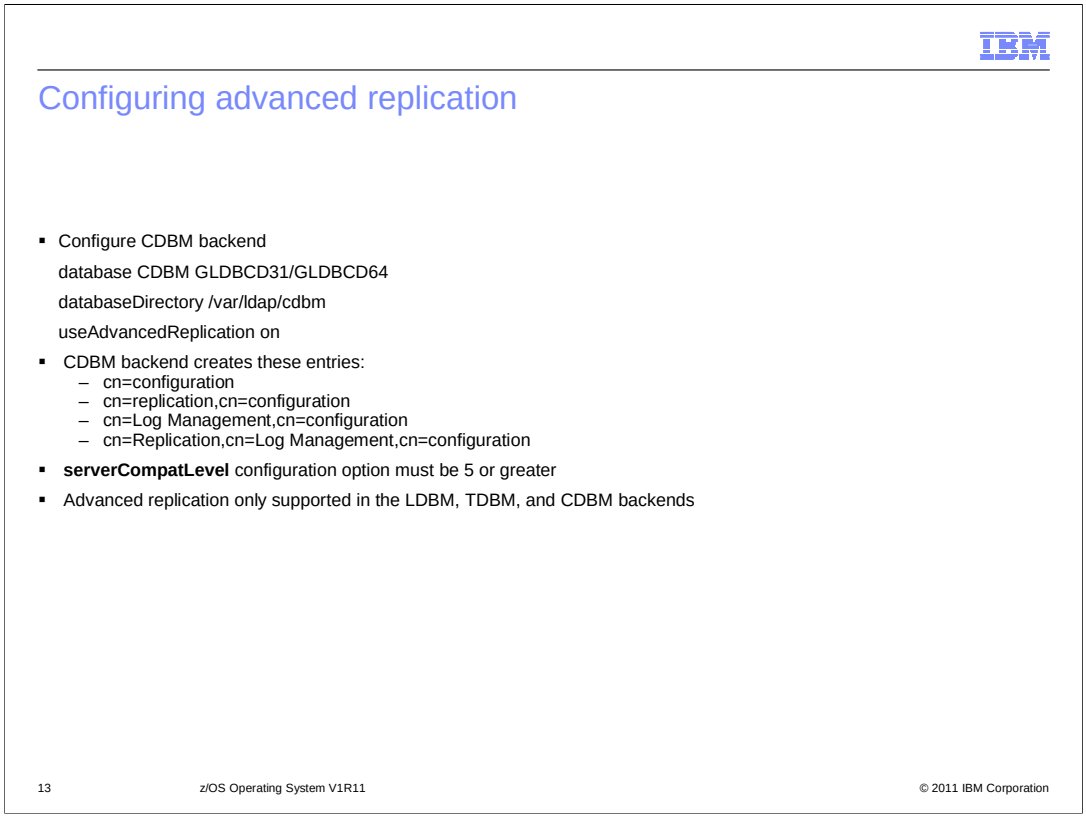

Starting in z/OS V1R11, the IBM Tivoli Directory Server for z/OS introduced a file-based configuration backend called CDBM. The CDBM backend must be configured to use advanced replication since the backend contains advanced replication configuration entries. The **serverCompatLevel** configuration option must be 5 or greater when the CDBM backend is configured. Also, the **useAdvancedReplication** configuration option in the CDBM backend section must be set to **on** so that advanced replication entries are allowed to be added to the server.

The CDBM backend has two suffixes **cn=configuration** and **cn=ibmpolicies.** The CDBM entries that are directly related to advanced replication configuration are: cn=configuration, cn=replication,cn=configuration, cn=Log Management,cn=configuration, and cn=Replication,cn=Log Management,cn=configuration.

Advanced replication is only supported in the LDBM, TDBM, and CDBM backends. Entries must be added or modified in these backends to successfully configure advanced replication. In order to replicate the schema entry, the cn=ibmpolicies entry must be modified to be a replication context.

An example of configuring a master-replica topology

- Modify cn=configuration entry on master (supplier) server

  dn: cn=configuration

  changetype: modify

  replace: ibm-slapdserverid

  ibm-slapdserverid: Master
- Modify cn=configuration entry on replica (consumer) server

  dn: cn=configuration

  changetype: modify

  replace: ibm-slapdserverid

  ibm-slapdserverid: Replica

The following slides explain how to configure an advanced replication master-replica topology. In this example, the master server is running on host server1.us.ibm.com and non-secure port 389; while the replica server is running on host server2.us.ibm.com and non-secure port 389.

Before advanced replication topology entries are added to the master server, the **ibm-slapdserverID** attribute value in the **cn=configuration** entry should be modified on both servers. By default, the **ibm-slapdserverID** attribute value is a randomly generated uuid value, which is similar to the **ibm-entryuuid** attribute values that are automatically generated when an entry is added to the z/OS LDAP server.

When bound as the LDAP administrator, the **ldapmodify** utility can be used to modify the **cn=configuration** entries on both servers. The **ibm-slapdserverId** attribute value in the **cn=configuration** entry on the master server is changed to “Master” while on the replica server it is changed to “Replica”. By modifying the **ibm-slapdserverId** attribute values in the cn=configuration entries, it makes it easier to configure advanced replication between the master and replica servers.

An example of configuring a master-replica topology (continued)

- Add the consumer server credentials entry to the replica (consumer) server:

```
dn: cn=Master server,cn=configuration
changetype: add
objectclass: ibm-slapdReplication
cn: master server
ibm-slapdMasterDN: cn=bindtoconsumer
ibm-slapdMasterPW: iamsupplier
ibm-slapdMasterReferral: ldap://server1.us.ibm.com:389
```


The consumer server credentials entry must be added to the replica or consumer server. This entry identifies the credentials that the master server will use to authenticate to the replica server. In this case, the supplier server authenticates to the consumer server with a bind DN of **cn=bindtoconsumer** and a password of **iamsupplier**.

If a client application attempts to update the replica server, it will receive a referral to the master server running on host server1.us.ibm.com and non-secure port 389.

An example of configuring a master-replica topology (continued) (1 of 5)

- Replication context entry:
  dn: o=ibm,c=us
  changetype: add
  objectclass: top
  objectclass: organization
  objectclass: ibm-replicationContext
  o: ibm
- Replica group entry:
  dn: ibm-replicaGroup=default, o=ibm,c=us
  changetype: add
  objectclass: top
  objectclass: ibm-replicaGroup
  ibm-replicaGroup: default

The following slides show creating an LDIF file called masterreplica.ldif. This LDIF file will contain the master-replica topology entries that will be added to the master server.

In this master-replica topology example, we want to replicate data under the o=ibm,c=us subtree. Therefore, the o=ibm,c=us entry must be designated as a replication context entry by adding the auxiliary objectclass **ibm-replicationContext.**

Directly under the replication context entry is the replica group entry. A replica group entry uses an auxiliary objectclass of **ibm-replicaGroup.** This objectclass also requires an **ibm-replicaGroup** attribute value.

An example of configuring a master-replica topology (continued) (2 of 5)

- Replica subentry to represent the master server:

```
dn: ibm-replicaServerId=Master,ibm-replicaGroup=default, o=ibm,c=us
changetype: add
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Master
ibm-replicationServerIsMaster: true
cn: Master
description: Master server of the topology
```


Directly under a replica group entry are replica subentries. Because this topology is using a master and a read-only replica server, a replica subentry is only needed for the master server. Read-only replicas do not need a replica subentry. A replica subentry has an auxiliary objectclass value of ibm-replicaSubentry and requires the cn, ibm-replicaServerId, and ibm-replicationServerIsMaster attribute types.

The replica subentry shown on this slide is for the master server. The ibm-replicaServerId attribute is set to Master which is the same value as the ibm-slapdServerId attribute value in the cn=configuration entry on the master server. Since the ibm-replicationServerIsMaster attribute value is set to true, it indicates that this server is a master for the o=ibm,c=us replication context. If the ibm-replicationServerIsMaster attribute value is set to false in this entry, this server would be a forwarder or cascading server.

An example of configuring a master-replica topology (continued) (3 of 5)

- Supplier server credentials entry:

```
dn: cn=ReplicaBindCredentials, o=ibm,c=us
changetype: add
objectclass: ibm-replicationCredentialsSimple
cn: ReplicaBindCredentials
replicaBindDN: cn=bindtoconsumer
replicaCredentials: iamsupplier
description: Bind Credentials on master to bind to replica
```


The supplier server credentials entry is used to specify the credentials that the master server will use to authenticate with the replica server. Since this supplier server credentials entry has an objectclass value of **ibm-replicationCredentialsSimple**, the master server will use a simple bind to the replica server with a bind DN of **cn=bindtoconsumer** and a password of **iamsupplier**. The bind credentials specified in this supplier server credentials entry are the same credentials specified in the consumer server credentials entry, cn=Master server,cn=configuration.

Advanced replication also supports authenticating to the consumer server using a SASL EXTERNAL bind by using a supplier server credentials entry that has an objectclass value of **ibm-replicationCredentialsExternal**.

An example of configuring a master-replica topology (continued) (4 of 5)

- Replication agreement entry from the master to the replica:

```
dn: cn=Replica, ibm-replicaServerId=Master,ibm-replicaGroup=default,     o=ibm,c=us
changetype: add
objectclass: top
objectclass: ibm-replicationAgreement
cn: Replica
ibm-replicaConsumerId: Replica
ibm-replicaUrl: ldap://server2.us.ibm.com:389
ibm-replicaCredentialsDN: cn=ReplicaBindCredentials, o=ibm,c=us
description: Replication agreement from master to replica
```


Replication agreement entries use a structural objectclass value of **ibm-replicationAgreement** and require the **cn**, **ibm-replicaConsumerId**, and **ibm-replicaUrl** attribute types.

Since the replication agreement entry on this slide is under the **ibm-replicaServerId=Master,ibm-replicaGroup=default,o=ibm,c=us** subtree, this entry identifies the replication path from the master to the replica. The replica server resides on host server2.us.ibm.com and non-secure port 389 and has a server ID of Replica. A SSL connection to the replica can be used by specifying ldaps:// as the prefix on the **ibm-replicaURL** attribute value and specifying the necessary LDAP SSL configuration options in the server configuration file.

The **ibm-replicaConsumerId** attribute value must be the same as the **ibm-slapdServerId** attribute value in the cn=configuration entry on the replica server.

The credentials that the master server uses to authenticate with the replica server are contained in the cn=ReplicaBindCredentials,o=ibm,c=us entry.

An example of configuring a master-replica topology (continued) (5 of 5)

- Put the master server into maintenance mode:
  - F *LDAPSRV,*MAINTMODE ON
- Add the entries to the master server:
  - ldapadd -h server1.us.ibm.com -p 389 -D *adminDN* -w *adminPW* -f masterreplica.ldif –k –L
- Add the replica topology entries to the consumer server by using the **Replication topology** extended operation
  - ldapexop -h server1.us.ibm.com -p 389 -D *adminDN* -w *adminPW* -op repltopology -rc o=ibm,c=us
- Move the master server out of maintenance mode:
  - F *LDAPSRV*,MAINTMODE OFF

At this point, you now have the masterreplica.ldif file, which contains all of the replication topology entries. Before adding the replication topology entries to the master server, the master server should be put into maintenance mode.

Once in maintenance mode, the replication topology entries can be added to the master server using the **ldapadd** utility with the **–k** and **-L** options. The **-L** option sends the **Do not replicate** control to the master server so that the replication topology entries are not automatically replicated to the replica server. The **–k** option sends the **Server Administration** control to the master server so that the addition of entries continues even when the subtree becomes read-only because of a server ID mismatch.

Once the entries are added to the master server, the replication topology entries need to be added to the replica server. However, instead of manually adding the same entries to the replica server, the **Replication topology** extended operation will be used on the master server. This extended operation automatically propagates the replication topology entries to the replica server. The **ldapexop** utility can be used to perform the **Replication topology** extended operation.

Once the replication topology entries are on both servers, the master server should be moved out of maintenance mode and then the master-replica topology is ready. The master accepts updates on the o=ibm,c=us subtree and propagates them to the replica. When a client attempts to update the replica server, it returns a referral to the master. However, the replica does handle compare and search requests.

More information

- IBM Tivoli Directory Server for z/OS Administration and Use manual (SC23-5191)

For more information on advanced replication, please see the IBM Tivoli Directory Server for z/OS Administration and Use manual.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_V1R11_Config_Adv_Rep_TivoliDirectory.ppt

This module is also available in PDF format at: ../V1R11_Config_Adv_Rep_TivoliDirectory.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, disclaimer, and copyright information