Advanced Technical Skills (ATS) North America

Tivoli Storage Manager
Node Replication Technical Overview and Update

Dave Daun

IBM Advanced Technical Skills

djdaun@us.ibm.com

TSM Advanced Technical Support Team

- **Dave Canan**
  - ddcanan@us.ibm.com
- **Richard Crespo**
  - racrespo@us.ibm.com
- **Dave Daun**
  - djdaun@us.ibm.com
- **Tom Hepner**
  - hep@us.ibm.com

Topics

- **Overview**
- **Configuring for Replication**
- **Replication Best Practices**
- **Replication Rules**
- **Performing replication**
- **Querying replication results**

Overview

Node Replication

TSM Node Replication Overview

- ***Node replication* is the process of incrementally copying, or *replicating*, data belonging to backup-archive client nodes**
- **Data is replicated from one TSM server to another server**
  - The server from which client node data is replicated is called a *source replication server*
  - The server to which client node data is replicated is called a *target replication server*
  - A server can function as the source of replicated data for some client nodes and as the target of replicated data for other client nodes
- **The purpose of replication is to maintain the same level of files on the source and the target replication servers**
  - As part of replication processing, client node data that was deleted from the source replication server is also deleted from the target replication server
  - When client node data is replicated, only the data that is not on the target replication server is copied

TSM Node Replication Overview

- **If a disaster occurs and the source replication server is temporarily unavailable, client nodes can recover their data from the target replication server**
- **If the source replication server cannot be recovered, client nodes can be converted for backup operations on the target replication server**
- **The following types of client node data can be replicated:**
  - Active and inactive backup data together, or only active backup data
  - Archive data
  - Data that was migrated to a source replication server by Tivoli Storage Manager for Space Management clients
- **Only Tivoli Storage Manager V6.3+ servers can be used for node replication**
  - However, data for client nodes that are running a client level V6.3 or earlier can be replicated
  - Data that was stored on a TSM V6.2 or earlier server before it was upgraded to V6.3+ can be replicated

Advantages Of TSM Replication

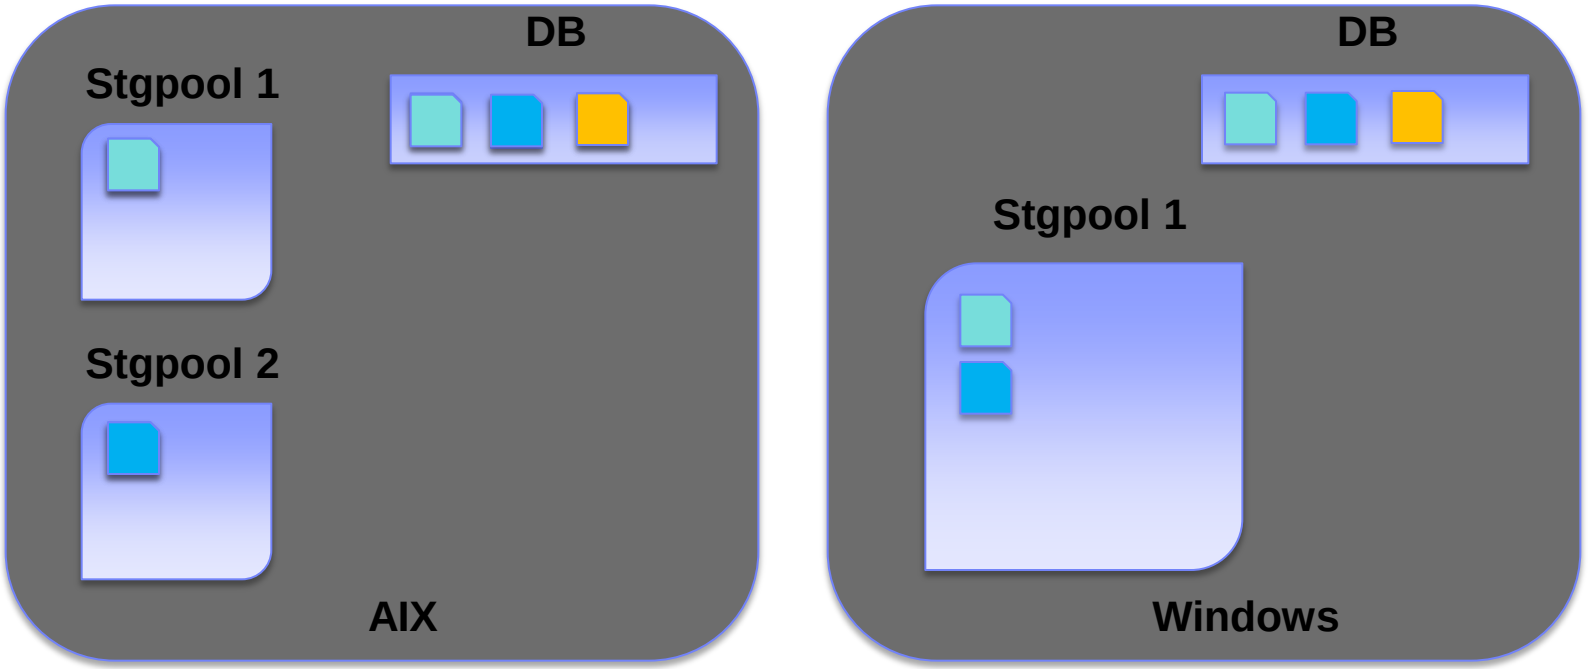

1. Different Operating Systems on source and target
2. Storage Pool layout can be different on each server
3. Meta-data and data always kept in sync
4. Data can be replicated from any storage pool
5. Objects without data are also replicated

TSM Node Replication Example

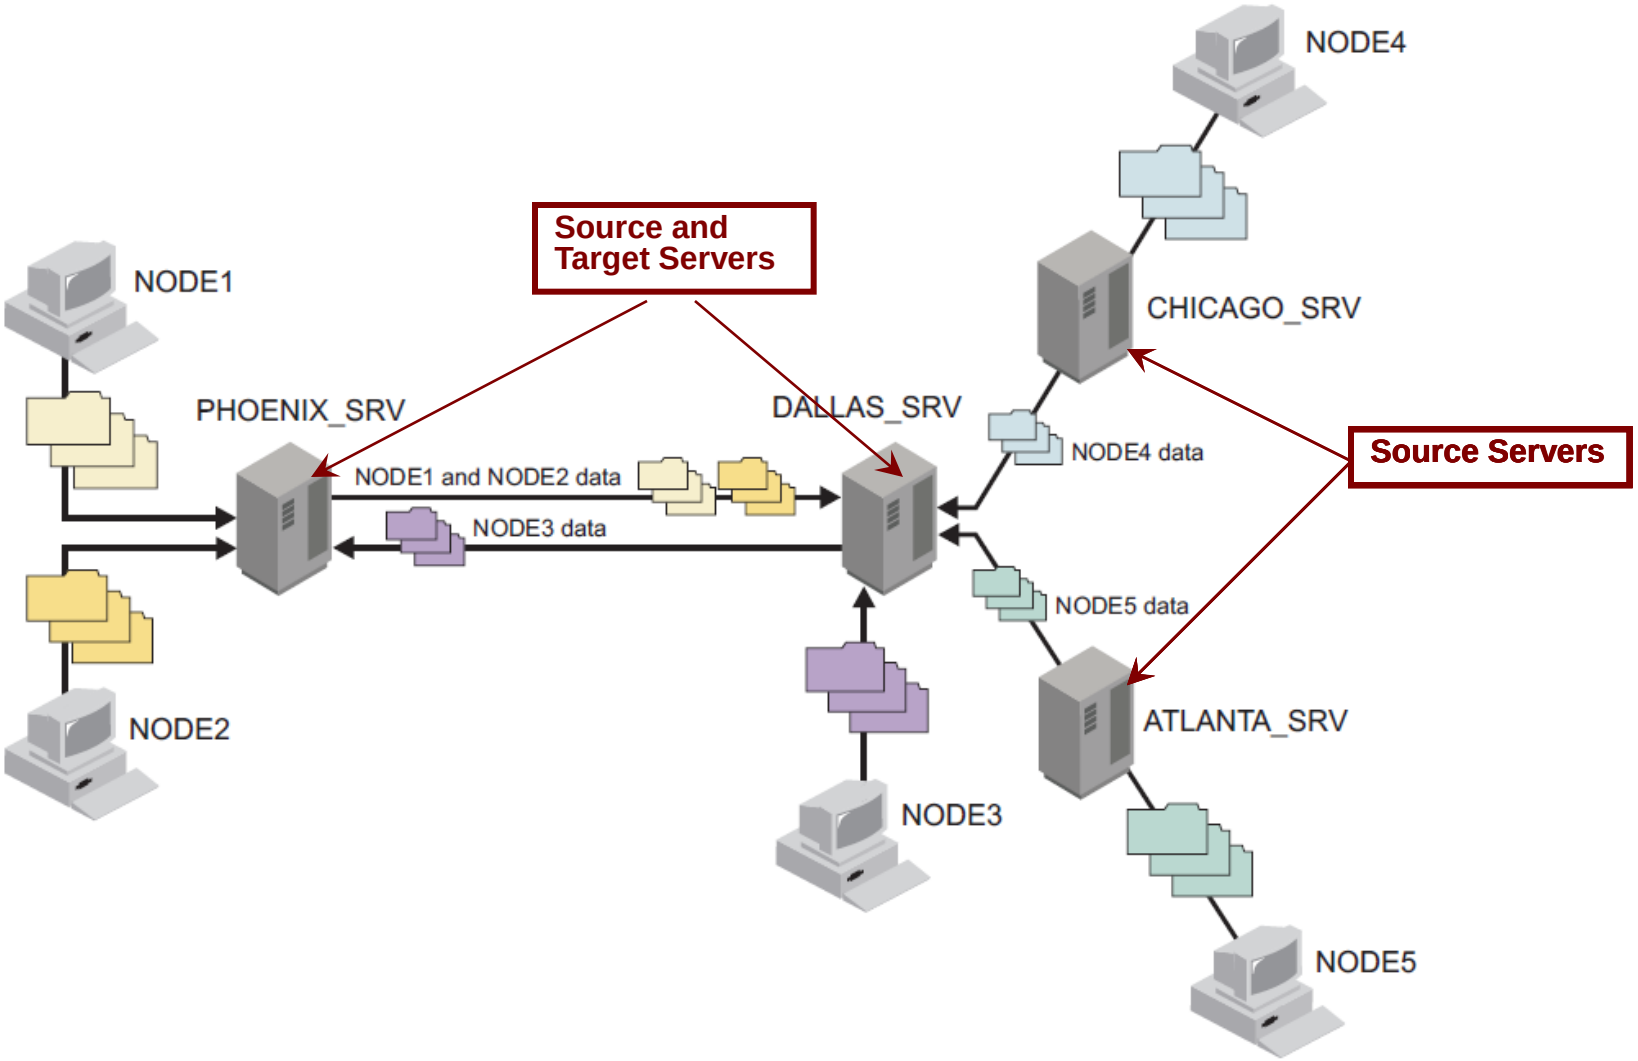

TSM Node Replication Considerations

- **Data that is replicated to a target server cannot be replicated again**
- **A source server can have only one replication target server per replication (i.e. it is not possible to replicate the same object to multiple TSM servers)**
- **SSL can be specified for secure communications between the servers**
- **Node Replication is De-duplication aware…**

TSM Node Replication and De-duplication

If the storage pool on the source server is...	*And the destination storage pool on the target server is...*	*The result is...*
Enabled for data de-duplication	**Enabled for data de-duplication**	**Only extents that are not stored in the destination storage pool on the target replication server are transferred**
Enabled for data de-duplication	**Not enabled for data de-duplication**	**Files are reassembled by the source server and replicated in their entirety to the destination storage pool**
Not enabled for data de-duplication	**Enabled for data de-duplication**	**Files are sent by the source server and replicated in their entirety.** **Server-side deduplication will occur on the target server**
Not enabled for data de-duplication	**Not enabled for data de-duplication**	**Files are replicated in their entirety to the destination storage pool**

Policy Management

- **When a client node is registered on a target server, the domain for the node is sent to the target server**
- **If the target server does not have a domain with the same name the node on the target server is placed in the default domain on the target server and replicated objects will be bound to the default management class**
- **To maintain the same number of file versions on the source and the target servers, the source server manages file expiration and deletion**
- **If a file on the source replication server is marked for deletion, the source replication server directs the target replication server to delete the file immediately**
- **Expiration processing on the target replication server is disabled for replicated data**

Policy Management

- **If a client node is removed from replication on the target server, the policies on the target server are enabled**
  - Data on the target server is then managed by the policies that are on the target replication server, and expiration processing can delete expired files
- **Policies that are defined on replication servers and are dissimilar can cause undesirable side-effects**
  - If a node is removed replication on the target server and the policy for that node on the target server specifies fewer versions than the source server policy did, then the extra versions on the target server will be deleted by inventory expiration
- **Configure the Domains, management classes and storage pools similarly on the source and target servers**
  - Use Enterprise Configuration
- **If a replicated file is deleted on the target server, it is eligible for replication then the next time that replication runs on the source server for the client node that owns the file**

Target Management Class Fields

Attributes Honored

Management Class Attributes

Default Mgmt Class
Migration Destination

Copygroup Attributes

Storage Pool Destination
TOC Destination

Attributes Ignored

Management Class Attributes

Space Management Technique
Auto-Migrate on Non-Use
Migration Requires Backup

Copygroup Attributes

Versions Data Exists
Versions Data Deleted
Retain Extra Versions
Retail Only Version
Copy Mode
Copy Serialization
Copy Frequency

Replication Rules

- **Replication rules control what data is replicated and the order in which it is replicated**
- **Replication rules are assigned separately to each data type:**
  - backup data
  - archive data
  - space-managed
- **Replication rules are applied collectively at the server level to all client nodes that are configured for replication, at the individual client node level and at the file space level**
  - File space rules take precedence over rules for individual nodes
  - Rules for individual nodes take precedence over server rules

Replication Rule Hierarchy

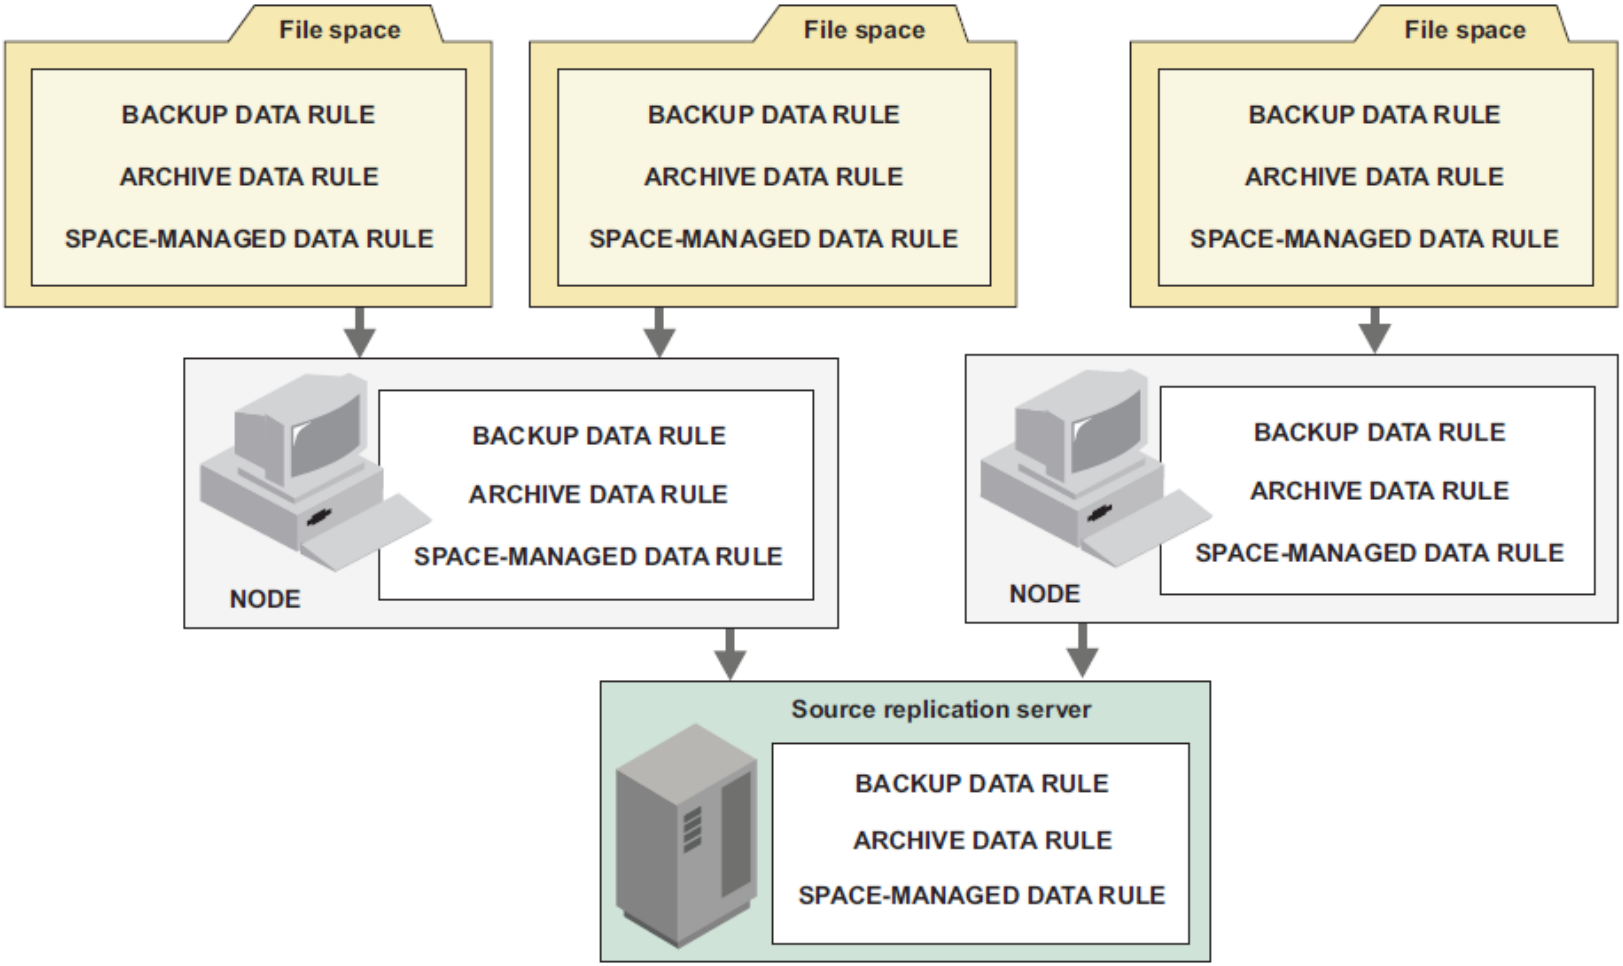

Replication Rules

- **The TSM server uses these predefined rules:**
  - ALL_DATA – replicates active and in-active versions with normal priority
  - ACTIVE_DATA -- replicates only active versions with normal priority
  - ALL_DATA_HIGH_PRIORITY -- replicates active and in-active versions with high priority
  - ACTIVE_DATA_HIGH_PRIORITY -- replicates only active versions with high priority
  - DEFAULT – Replicates data according to the rule that is assigned to the data type at the next higher level in the replication-rule hierarchy
  - NONE – Data is not replicated
- **The entire set of replication rules is predefined with these values:**
  - All file space and client node level rules are set to DEFAULT
  - All server level rules are set to ALL_DATA

Replication States

- **Replication state indicates whether replication is enabled or disabled**
- **Replication states apply to these objects:**
  - Replication Rules -- If a replication rule is disabled, data to which the rule is assigned is not replicated
  - Individual Client Nodes – if client node replication is disabled, backup, archive, and space-managed data in file spaces that belong to the client node is skipped during replication processing
  - Data Types in File Spaces – if replication is disabled for one or more data types in a file space, then files with that data type (backup, archive, space management) is skipped during archive processing

Running the Replication

- **Replication can be run for data that belongs to one or more client nodes or the data that belongs to a defined group of client nodes**
- **Start replication with the REPLICATE NODE administrative command on the source server**
  - It can be run manually
  - It can be run on a regular basis as part of an administrative schedule
- **The REPLICATE NODE command starts a single process**
  - Can have multiple data sessions flowing to the target server controlled by the MAXSESSIONS parameter

The REPLICATE NODE Command

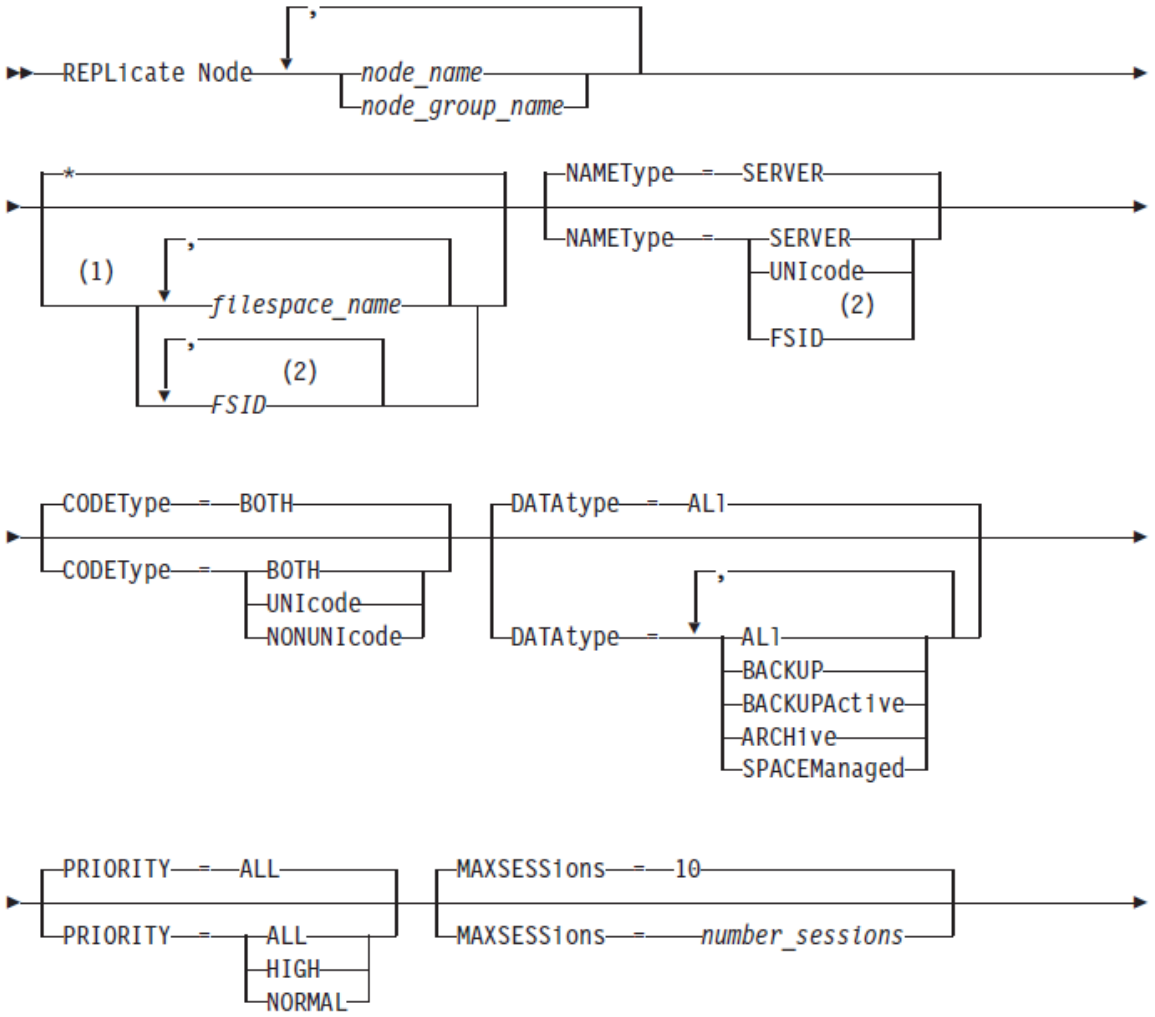

Monitoring Replication

- **Command: QUERY REPLNODE ***
  - Lists the number of files for a given set of nodes on both the source server and the target server
  - Use to see if a node is fully replicated
- **Command: QUERY REPLICATION node_name**
  - List the status of running or completed replication events over some period of time
  - Can provide detailed info about the type and number of objects replicated
- **Tivoli Monitoring for TSM (formerly TSM Monitorng and Reporting)**
  - New real-time and historical reports for Node Replication

TSM Node Replication After a Disaster

- **To restore, retrieve, or recall data from a target replication server, update the client options file to point to the target replication server**
  - Changes to node replication settings are not required
- **If a source replication server is unavailable for an extended period and backup operations are necessary client nodes on the target server can be converted to non-replicating nodes**
  - Non-replicating client nodes can back up, archive, or migrate data to a target replication server
  - Consider this move carefully as it may be resource intensive to synchronize back again to the source server
  - Client schedules are not replicated and will need to be defined on the target server

New Command: REMove REPLNode

- Removes a node from replication
- Replication configuration and object information in database is deleted for node
- Node is updated to a non-replicating node
- Data is not deleted (source and target)
- Must be issued on all servers that have the node configured for replication

```
>> --REMove REPLNode-----+--node_group_name--+-+---->
                         ^   '--node_name-----' |
                         ^---------,-----------'
```

Configuring for Replication

Node Replication

Node Replication Overview

Source TSM Server

1 Enable source server to target server communication

2 Set target server as default target for replication

3 Enable nodes to replicate and modify default and filespace replication rules if needed

4 Set storage hierarchy , domain policies

5 At proper time, preview and initiate replication for desired nodes and data (REPLICATE NODE)

Target TSM Server

7 DR: restore node data

6 TSM administrator views replication results

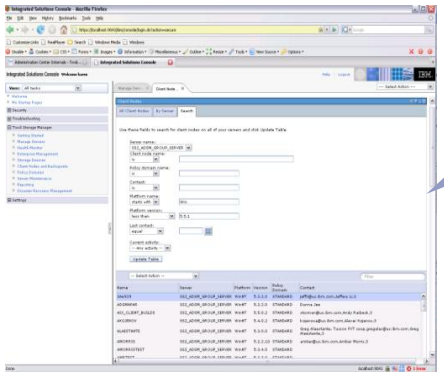

Admin Center & TCR

Initial configuration

Other

General Replication Steps

1. Create the server definitions on the source and target servers
2. Set the default target replication server
3. Setup policy and storage hierarchy on the target server
4. Determine nodes and filespace data to replicate
5. Configure nodes for replication
6. Modify filespace, node or server replication rules if needed
7. Validate replication rules (VALIDATE REPLICATION command)
8. Perform replication (REPLICATE NODE command)
9. Verify results (QUERY REPLICATION command)
10. Monitor replication (QUERY REPLNODE command)

1. Create server definitions
2. Issue SET REPLSERVER
3. Set policy on target server
4. Determine nodes to replicate

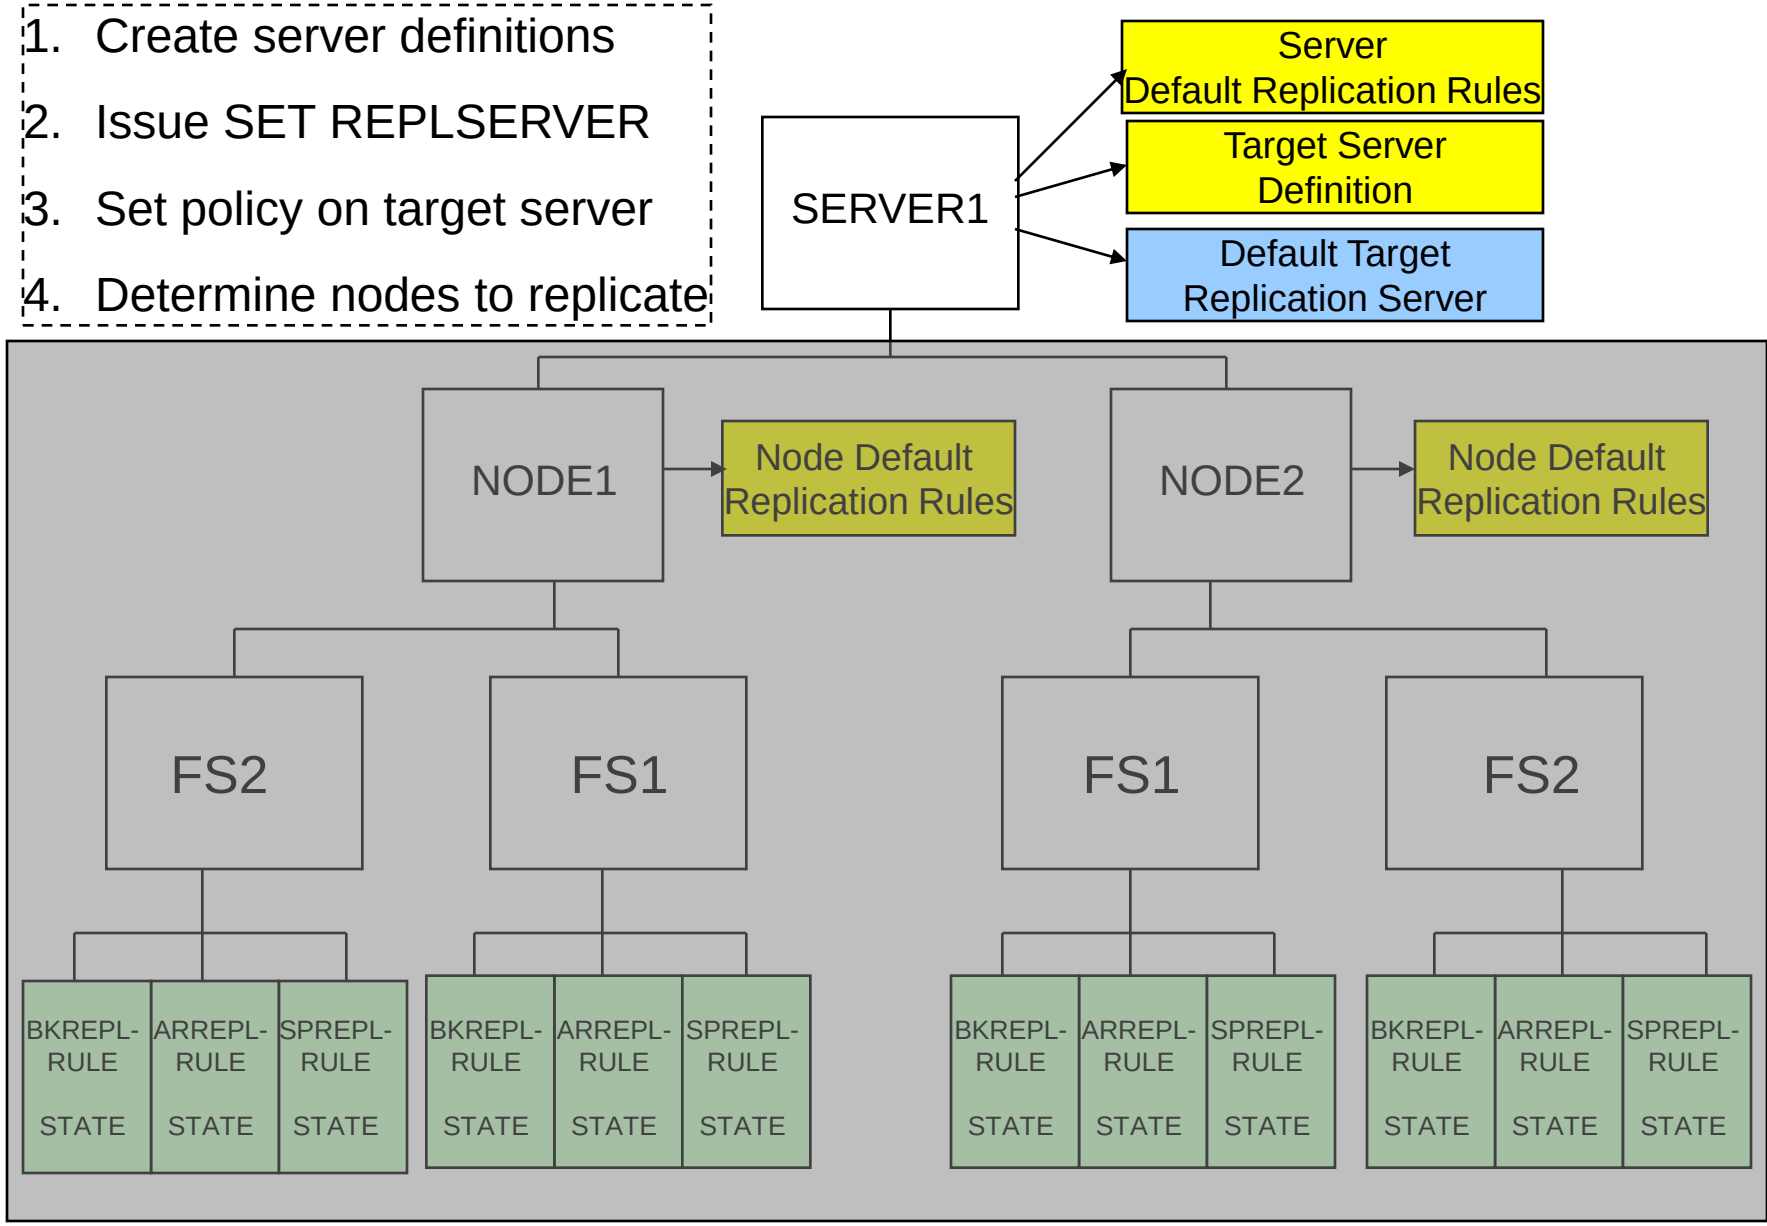

Server to Server Communication

- **Define the source and the target replication servers to each other, including Secure Sockets Layer (SSL) communications, if required**
  - *Target Server*: SET SERVERNAME, SERVERPASSWORD, SERVERHLADDRESS, SERVERLLADDRESS
  - *Target Server*: SET CROSSDEFINE ON
  - *Source Server*: SET SERVERNAME, SERVERPASSWORD, SERVERHLADDRESS, SERVERLLADDRESS
  - *Source Server*: DEFINE SERVER (using correct servername, serverpassword, serverhladdress, serverlladdress)

Set the Default Target Replication Server

- ***Source server*: SET REPLSERVER target_server_definition_name**
  - Defines the target replication server to be used for replication
  - Stored in system parameters--use QUERY STATUS to view the default target replication server name
  - Target server is checked during replication to ensure that it is a TSM 6.3+ server
  - Four server replication rules are created when TSM V6.3+ is installed that are used to determine what data in each filespace is eligible to be replicated

Managing Policy Across the Servers

- **While it is not required, it is recommended and considered best practice that the policies on the source and target servers are configured the same**
  - Enterprise configuration can be used to keep the policies synchronized
- **During replication the source server sends the management class name to the target server**
  - If there is a matching management class, the data is bound to that management class
  - If the target server does not have the management class defined, the data is bound to the default management class
- **Version control and retention is managed by the source server.**
- **Policies on the target server (think expiration) will only be used to manage replicated client data should the node be removed from replication on the target server**

Determine Which Nodes and Filespaces to Replicate

- **Which nodes have data that needs to be replicated?**
- **Which filespaces belonging to a node have data that needs to be replicated?**
- **What type of data in the filespaces should be replicated (backup, archive, space managed)?**
  - If backup data is to be replicated, should only the active data be replicated?
- **Is any of the data in the filespaces more important than others and needs to be replicated first?**

5. Configure nodes for replication
6. Modify replication rules if desired

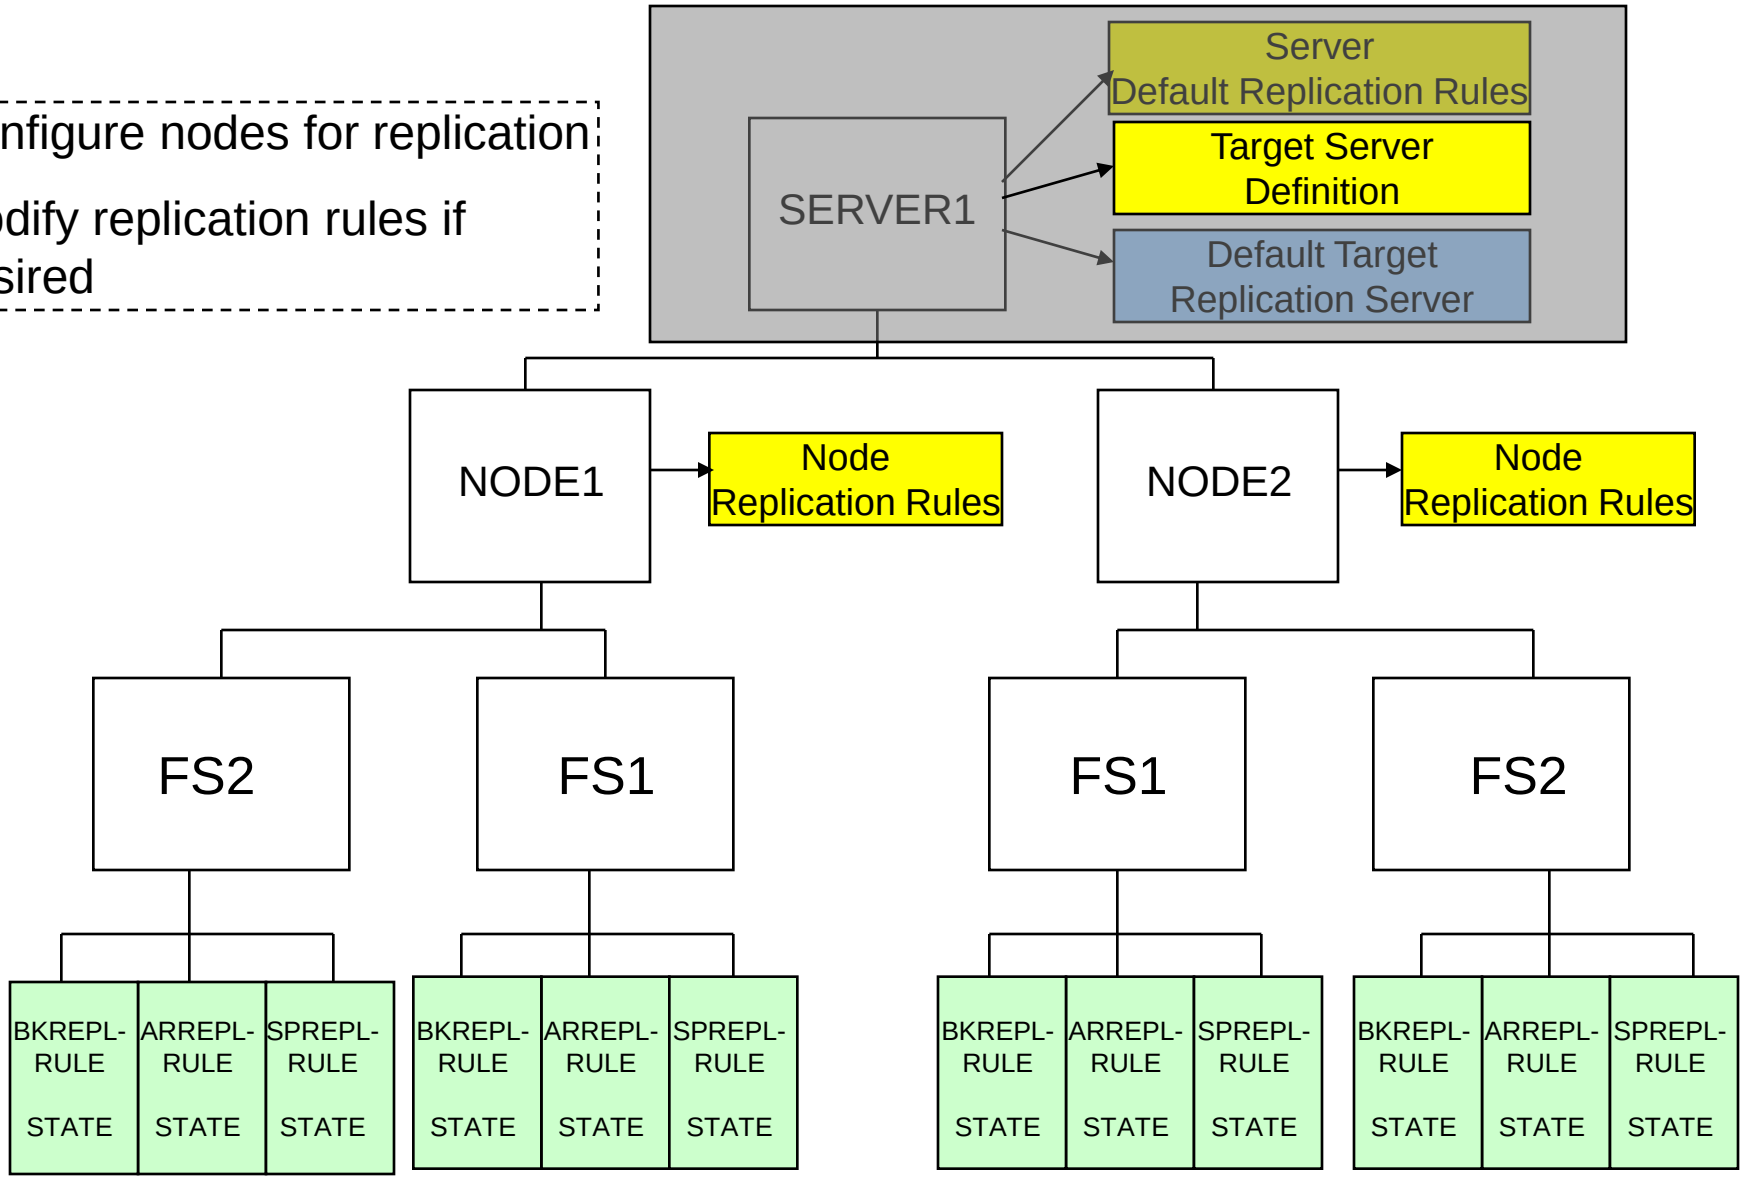

Configure Nodes for Replication

- **UPDATE / REGISTER NODE node_name REPLSTATE=state REPLMODE=mode**
  - Specifies the replication state of the node and is used to determine if replication should be run for the node
  - Specifying this parameter will also configure a node for replication if the node is not currently configured
  - A node whose definition has REPLSTATE=NONE and REPLMODE=NONE is not configured for replication
- **REPLSTATE = Enabled|DISabled|None**
  - Enabled--Specifies that node is configured and enabled for replication.
  - DISabled--specifies that node is configured for replication but it is temporarily disabled and replication will skip the node when in this state.
- **REPLMode = SYNCSEnd | SYNCRECeive**
  - Specifies the replication mode of the node is being set in order to synchronize exported and imported data on two servers
  - SYNCSEnd--Specifies that during replication, data belonging to this node is synchronized with data on a target server. When the synchronization is completed, the REPLMODE parameter for this node on the source server is set to SEND. This value should be specified on the source server.
  - SYNCRECeive--Specifies that during replication, data belonging to this node is synchronized with data on the source server. When the synchronization is completed, the REPLMODE parameter for this node on the target server is set to RECEIVE. This value should be specified on the target server.

Configure Nodes for Replication

- **When a node is initially defined, by default it is set to REPLMODE of NONE**

- **When a node is initially defined or updated to REPLSTATE of ENABLED, the:**

  - Filespace level replication rules are set to DEFAULT
  - Node level replication rules are set to DEFAULT
  - This means that data will be replicated based on what has been specified at the Server level for the replication rules.

Initial Configuration of a Node for Replication

- **Method #1 (Node does not exist yet on Target Server.)**
  - Use the register/update Node commands on source server, specifying REPLSTATE parameter but not the REPLMODE parameter
  - REPLMODE automatically set to SEND on source server node
  - When replication first occurs, a node will be defined on the target server with REPLSTATE set to RECEIVE
- **Method #2 (Node does exist on Target Server and Import/Export has been used to import some data to the target server)**
  - On source server, issue update node command with REPLSTATE=ENABLED and REPLMODE=SYNCSEND
  - On target server, issue update node command with REPLSTATE=ENABLED and REPLMODE=SYNCRECEIVE
  - After first replication occurs, REPLMODE will be set to SEND on source server and REPLMODE will be set to RECEIVE on target server.

Populating the Target Server

- **Two basic methods to populate target server**
  - Method 1 – Replicate from scratch
    - Best if source and target are in close proximity
    - All eligible data is sent
  - Method 2 – Synchronize and Replicate
    - Best if source and target are at great distance or if bandwidth is limited
    - Use media-based Export/Import to populate target server
    - Replication with SYNC links the source and target objects

Sync Example

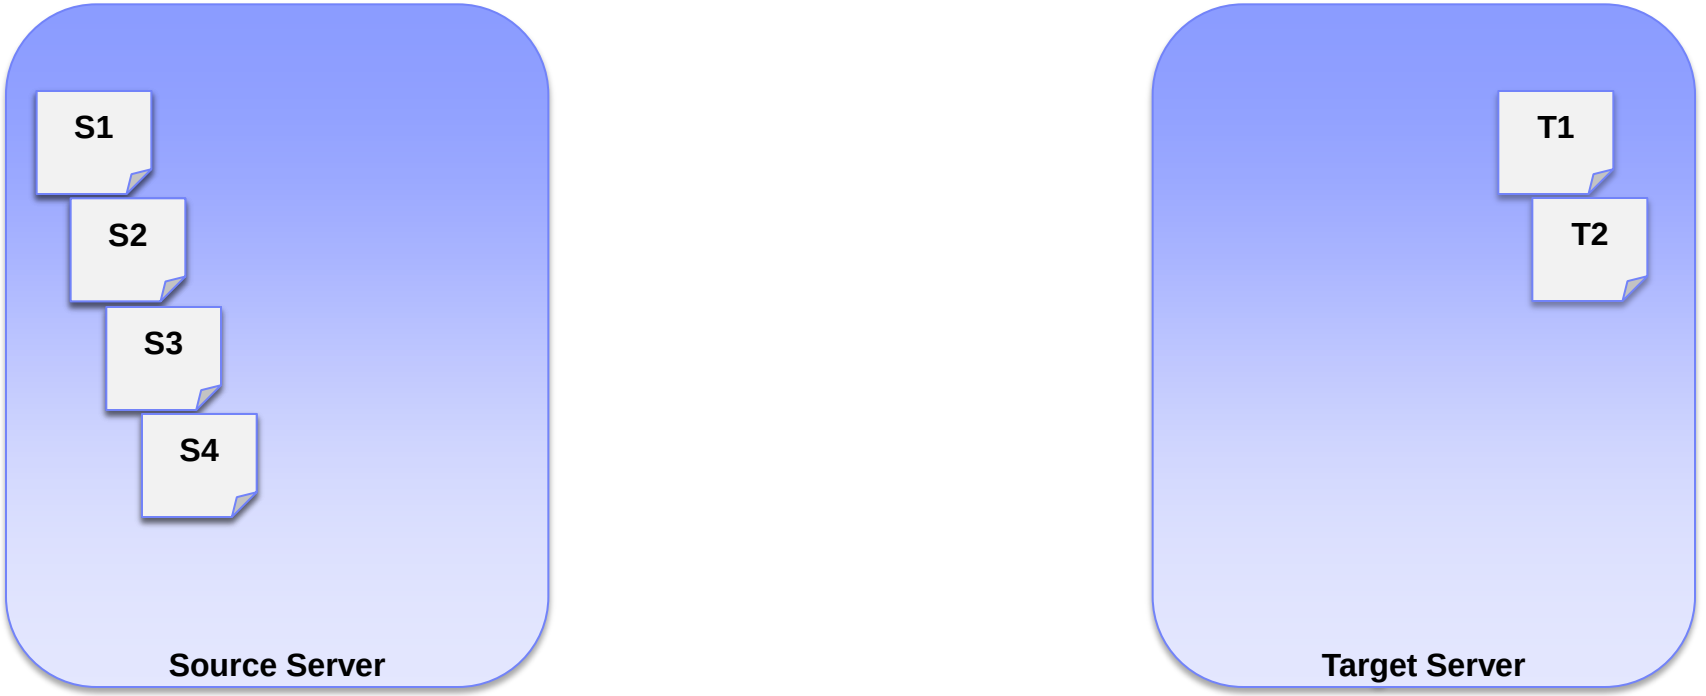

Sync Example

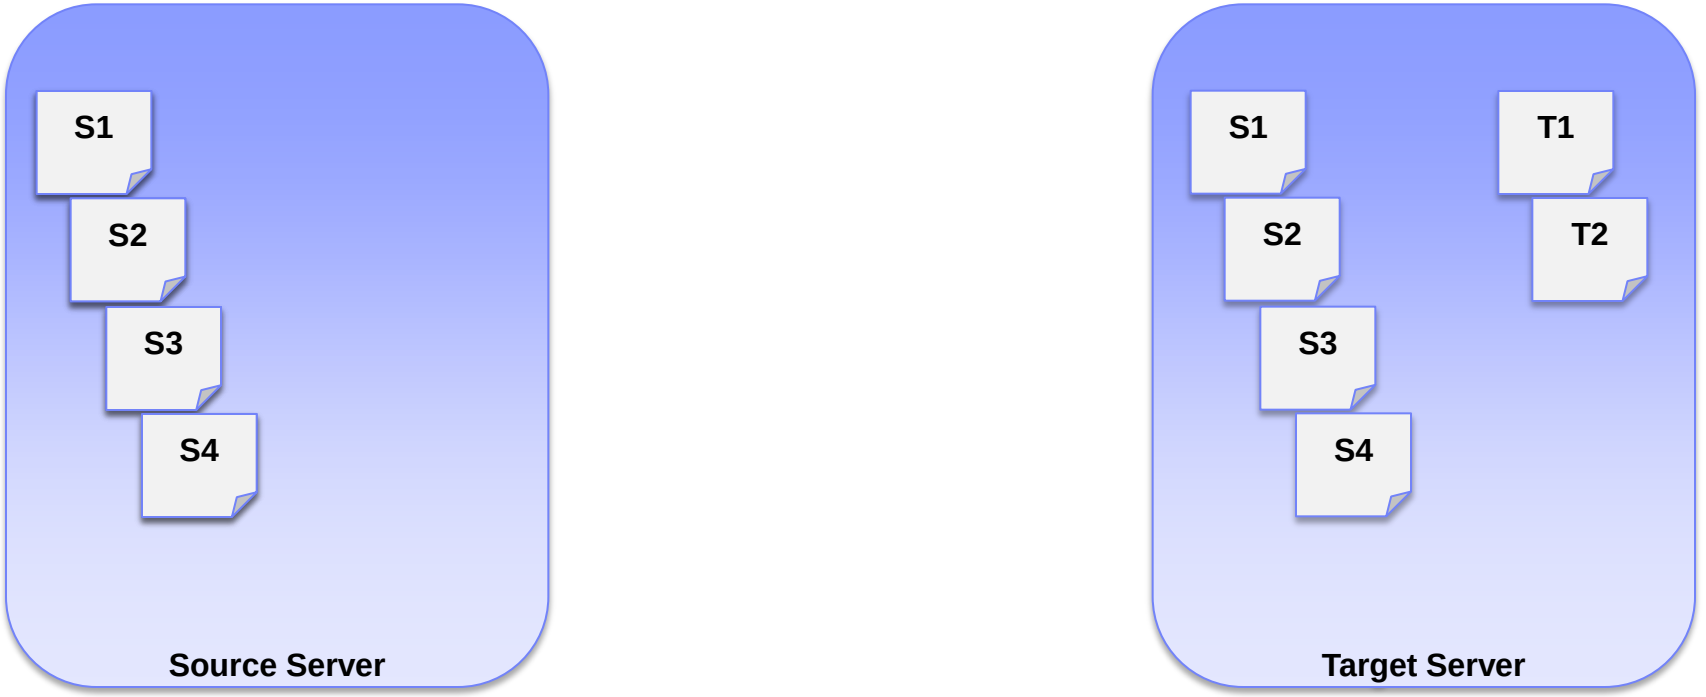

1. **Export/Import used to transfer data to target server**

Sync Example

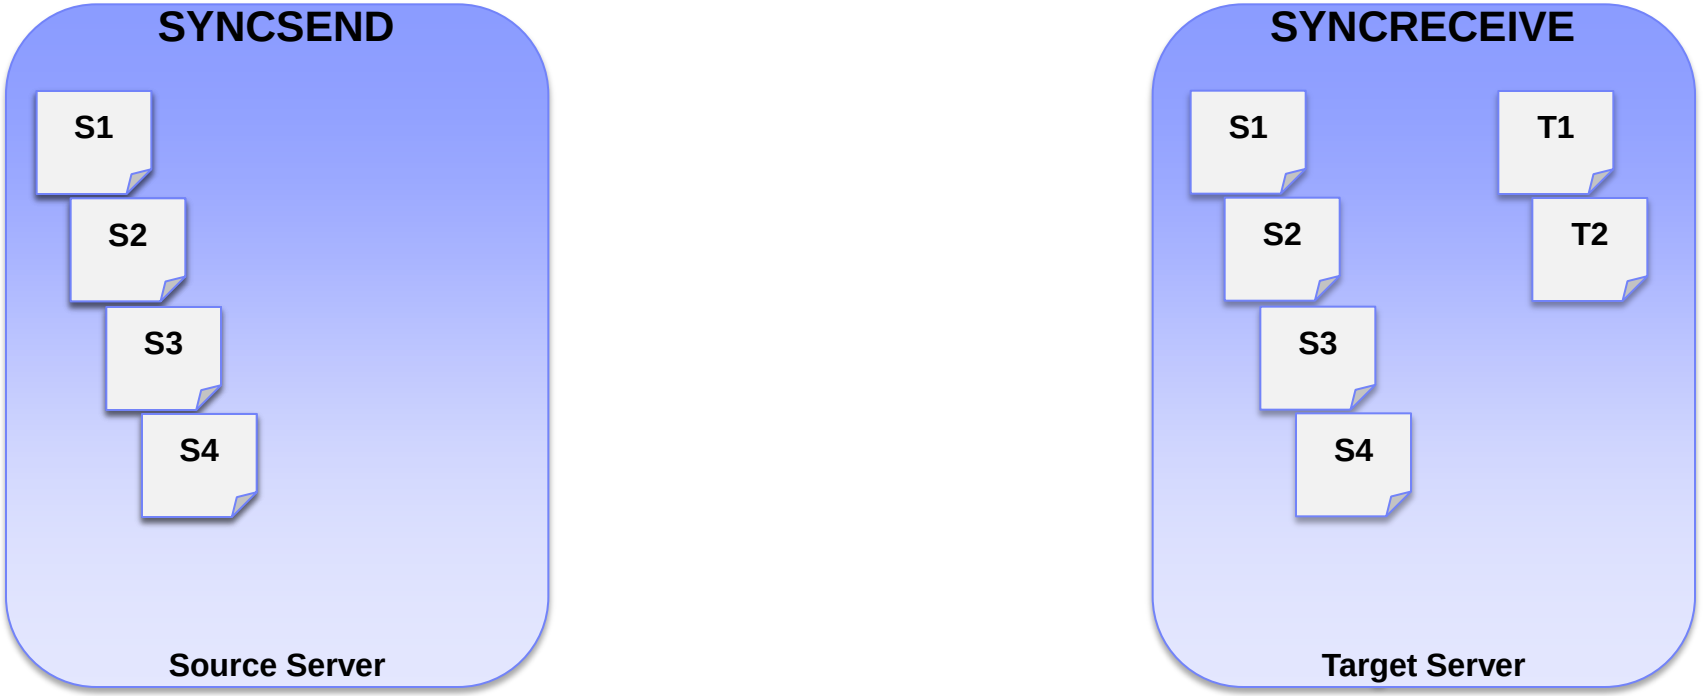

1. **Export/Import used to transfer data to target server**
2. **REPLMODE set appropriately on each server**

Sync Example

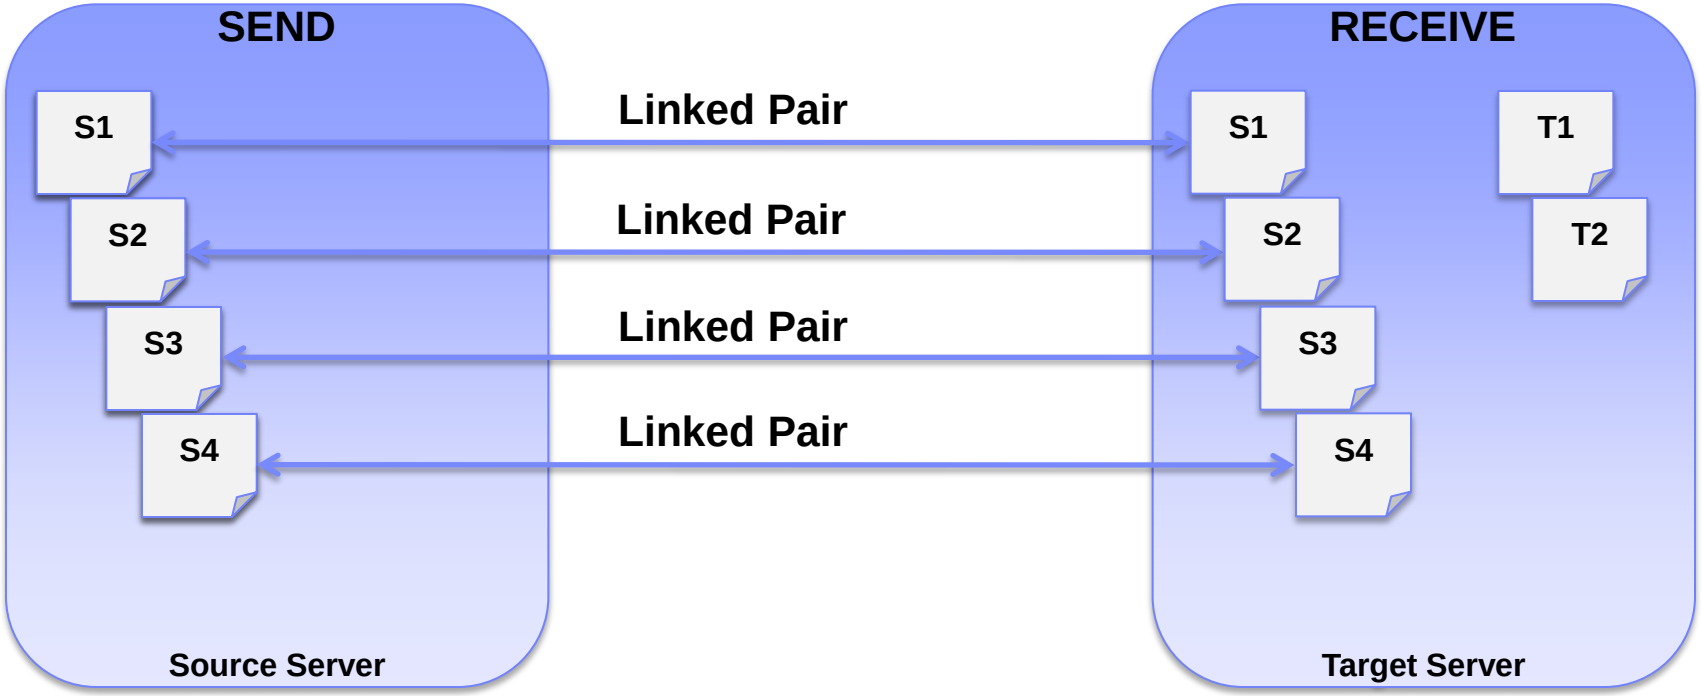

1. Export/Import used to transfer data to target server
2. REPLMODE set appropriately on each server
3. Replication will link the common objects and change the REPLMODE

Sync Example

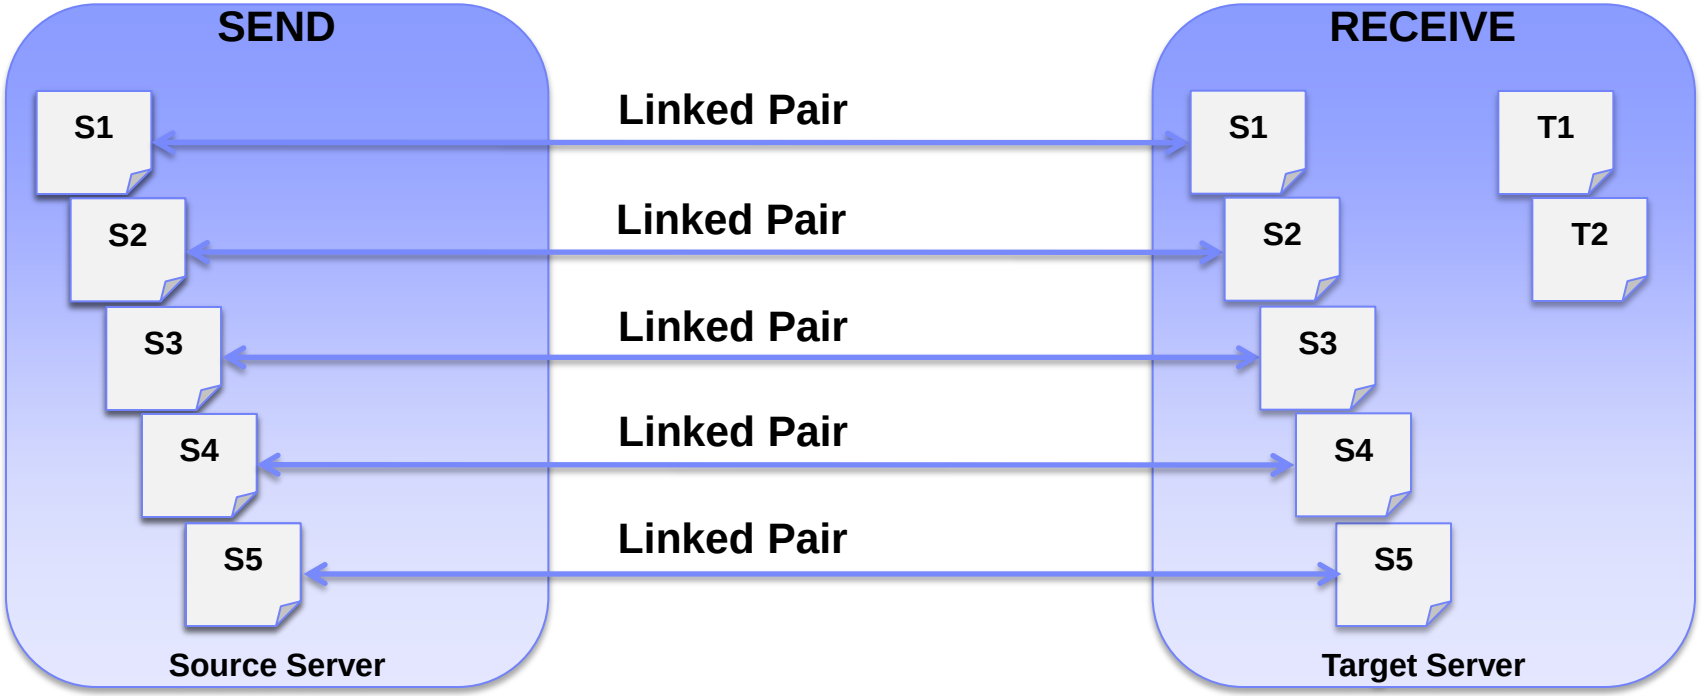

1. Export/Import used to transfer data to target server
2. REPLMODE set appropriately on each server
3. Replication will linked the common objects and change the REPLMODE
4. Subsequent replications will copy new objects

Replication Best Practices

Node Replication

Read the TSM Wiki Item on Replication

- *Guidelines for Replication*
- https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Storage%20Manager/page/Guidelines%20for%20node%20replication

1. System Requirements
2. Network Bandwidth Considerations
3. Planning for Mount Point Consumption
4. TSM Server Level

Replication System Requirements

- Additional memory is required
  - above what is typically recommended for a TSM 6.3+ server.
  - running a server with insufficient memory may cause server performance to decrease or other operational issues
- Without deduplication
  - minimum of 64 GB of memory and 4 CPU cores
- With deduplication
  - minimum of 128 GB of memory and 8 CPU cores
- For larger workloads, you may need higher amounts of memory.

Replication System Requirements

- TSM Database Requirements
  - You will need additional TSM database space to track the files that have been replicated
- Estimate the replication database space
  - Issue QUERY OCCUPANCY for each node requiring replication
  - Capture the number of objects per node and the sum for all nodes
  - Multiply total objects times 300 to get total additional database bytes
  - Add another 10% for good measure
- Active Log
  - At least 64 GB for non-deduplicated workloads
  - 128 GB for deduplicated workloads
- Be sure the target side database is at least as big

Network Bandwidth Considerations

- Estimating the amount of data for replication
  - *For the initial replication:*
  - Use the QUERY OCCUPANCY command for each replication node—use the physical space values for non-deduplicated environments—use the logical space values when the target server is not deduplicated--use the logical space values and apply a deduplication ratio when both source and target are deduplicated
  - *For the subsequent incremental replications:*
  - Determine how data is backed up per day on each of the replicated nodes
  - Pick off the ANE4961I messages in the Activity Log
  - Use statistics from the Summary table

Network Bandwidth Considerations

- Test the bandwidth capacity
  - Select one or more nodes that have 500 GB to 1 TB of data--ensure that the data is typical of the data that you will replicate on a routine basis
  - Configure the nodes for replication
  - Select a window of time during which replication would normally run
  - If you plan to use SSL during replication, ensure SSL is enabled
  - Issue the REPLICATE NODE command to start the replication process
  - After the replication process completes, review the message ANR0327I in the Activity Log--use the values "Amount transferred" and "Elapsed time" to determine the number of bytes per hour that can be replicated
  - Run several tests and average the results

Network Bandwidth Considerations

- Determine the required replication time
  - Divide the replication throughput rate into the total number of bytes needed for replication
  - Do two calculations—one for initial replication and one for daily replication
- Determine the available replication window
  - Node replication, like other TSM processes, will not tolerate a significant amount of overlap with other server tasks especially Inventory Expiration
  - Do not run Node Replication during client backups
- What if the required replication time exceeds the replication window?
  - Tune the network and the replication process
  - Reduce the amount of replication data
  - Increase the network bandwidth

Replication Mount Point Usage

- The number of replication sessions can affect mount point usage
  - The default value for sessions is 10
  - More sessions can increase replication performance
- Deduplication can add to the number of mount points needed during replication
  - The NUMOPENVOLSALLOWED server option controls how many volumes are left open and mounted in a deduplication environment
  - Use this formula to determine the number of mount points needed in a deduplication environment: NUMOPENVOLSALLOWED * MAXSESSIONS
- Consider other workload (processes and sessions) which could consume mount points

TSM Server Level

- When using Node Replication, it makes sense to keep current on the TSM Server code level
  - TSM Server 6.3.3.100 includes some maintenance for Node Replication
- Keep in touch with the IBM Support Center for the latest information on Node Replication maintenance

Architecture: Nearby Replication

- **When source and target servers are close by**
  - Can utilize high bandwidth to transfer more data
  - Clients can access target server as easily as they access the source server
  - May be possible to replicate entire servers

- **Clustering vs Replication**
  - Clustering protects against HW failure
  - Replication protects against HW and media failure

Architecture: Long Distance Replication

- **When source and target are far apart**
  - How far is “far”? Farther than “near”!
  - Presumably at a completely different site—across town or across the country
  - This is what is often call “Electronic Vaulting”
- **Consider limiting to high-importance nodes**
  - May not want to replicate everything
  - OR consider replicating only active data
- **Reduced bandwidth requirement with deduplication**
  - Still need to have reasonable bandwidth
  - Target server must be configured to use sequential disk

Scheduling

- **If possible, allow sufficient time for IDENTIFY to process all data before replicating**
  - Allows replication to benefit from deduplication
- **Run storage pool reclamation after replication**
- **Run storage pool migration after replication**
  - Faster to replicate from disk than tape
- **Replication and storage pool backup can be run simultaneously**

Node Replication and Copy Storage Pools

- **Node Replication does not provide easy recovery from media failures on the source system**
- **It may be necessary to perform both Node Replication and Backup Storage Pool**
  - For media recovery purposes on the source server
- **It is possible to temporarily reverse the Node Replication but it is a manual, multi-step process that is a bit cumbersome and could be error-prone**
- **Make a reasonable assessment of the risks and costs when deciding whether or not to continue Backup Storage Pool**

Node Replication and Client Passwords

- **Node Replication will replicate client passwords to the target system**
- **For platforms other than Windows, this will allow clients that use PASSWORDACCESS GENERATE to attach directly to the target system and perform restores without having to manually provide the TSM password**
- **For Windows platforms:**
  - The TSM client on Windows stores the encrypted password associated with the TSM Server Name
  - Windows clients will prompt for the password when attempting to connect to the replication target server the first time
  - To prepare for disaster, all Windows clients should make a connection to the target server and provide the password manually—this will store the encrypted password on the client twice (once for the source server and once for the target server)

Replication Rules

Node Replication

What are Replication Rules?

- Replication rules are used to determine what files in a filespace are eligible for replication
- Replication process examines the rules to determine which files to replicate
- There are four attributes of a replication rule:
  - target replication server (default target replication server)
  - priority (HIGH or NORMAL)
  - replicate active-data only? (only applies to backup data)
  - state: enabled or disabled
- There are 6 replication rules
  - 4 server replication rules
  - DEFAULT - default replication rule that indicates to follow the replication rule hierarchy until a non-DEFAULT rule is specified.
  - NONE – replication is not performed, no data is replicated

Replication Rules Hierarchy

- Rules processed from filespace level up to server level
- If a rule at a filespace level for a given datatype has a value of DEFAULT, then node replication rule examined
- If a rule at a node level has a value of DEFAULT, then server replication rule examined
- When a non-DEFAULT rule encountered, search for a replication rule stops.
- A server level replication rule **cannot** have a value of DEFULAT

7. Issue Replicate node

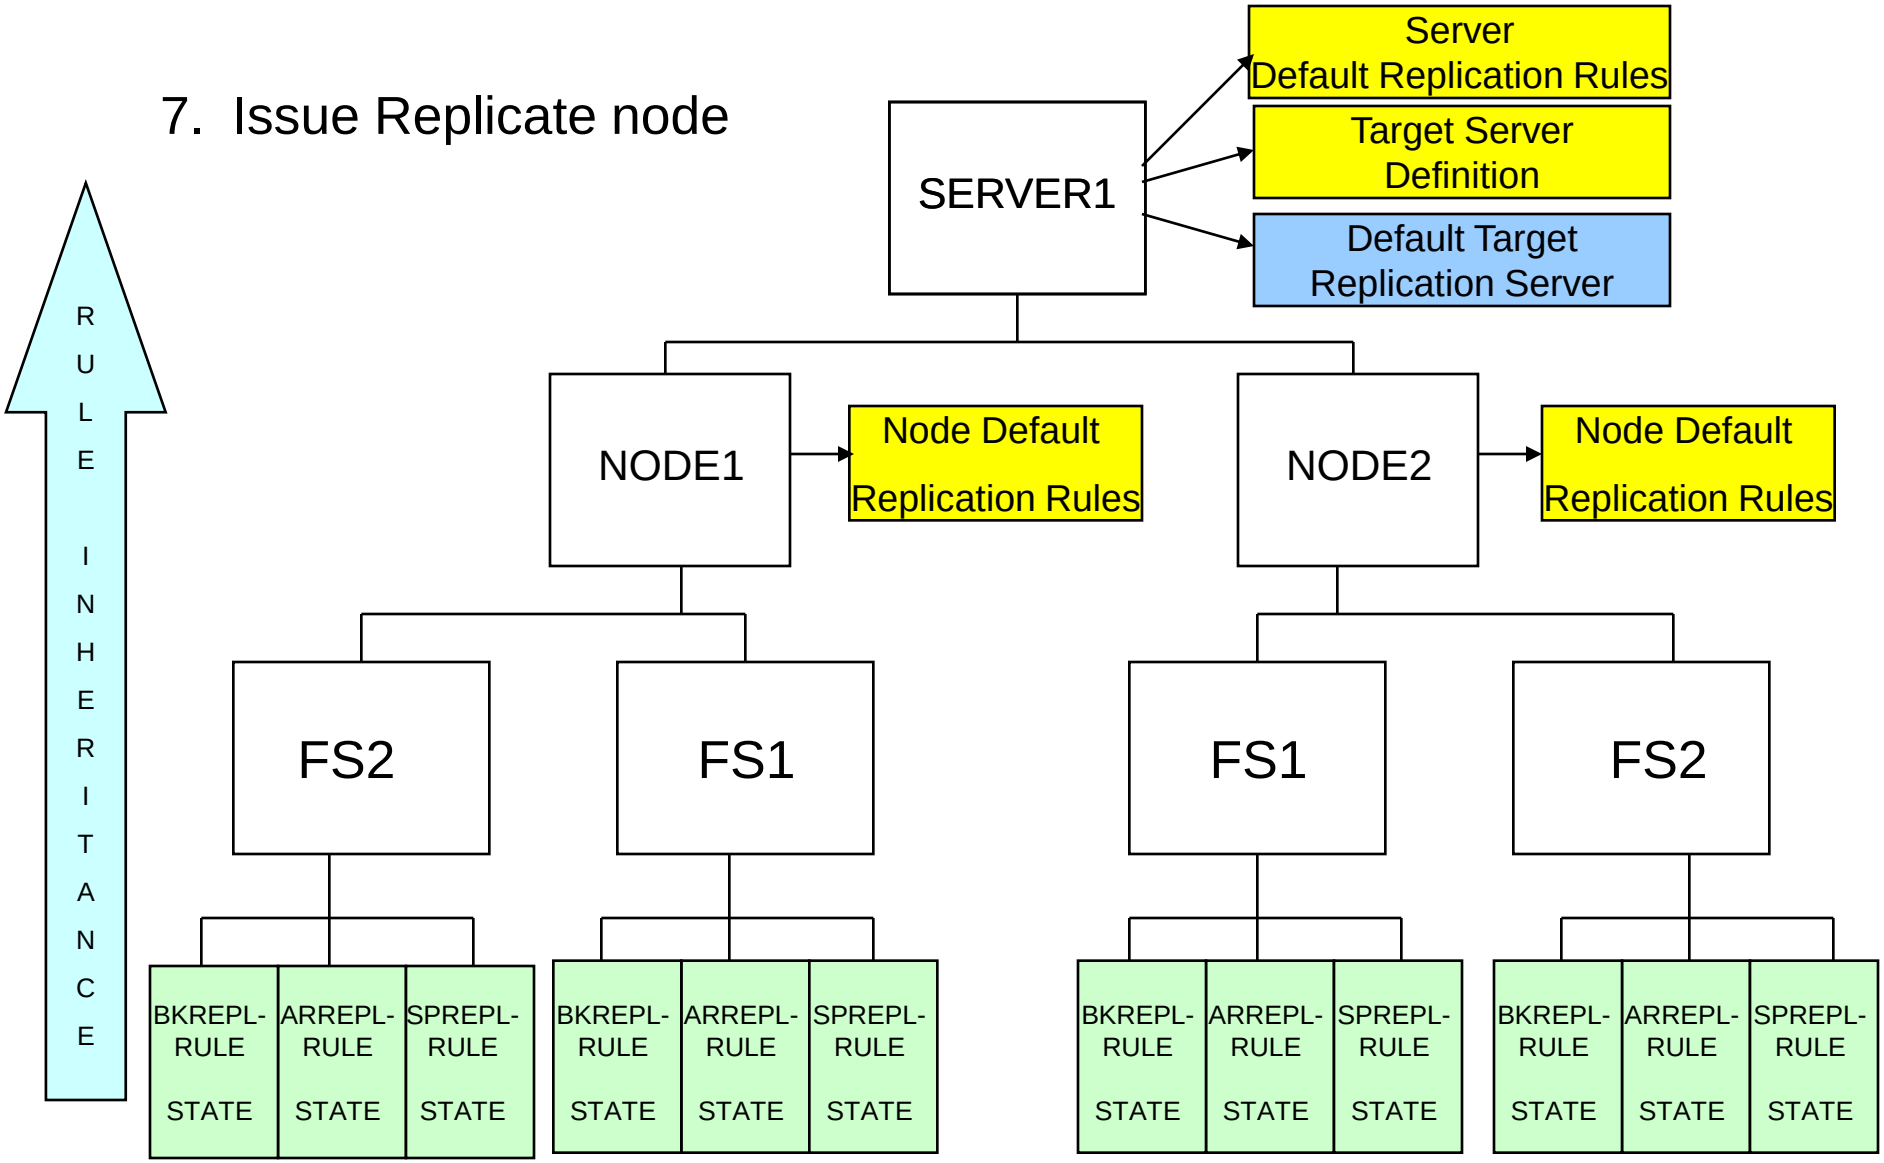

Four server replication rules (as initially defined)

ALL_DATA (all backup, archive or space managed data in the filespace)
- REPL_SERVER is set to the default replication server
- REPL_PRIORITY is set to NORMAL
- REPL_ACTIVE_DATA is set to FALSE.
- REPL_STATE is set to ENABLED (default)

ACTIVE_DATA (active backup data in filespace only)
- REPL_SERVER is set to the default replication server
- REPL_PRIORITY is set to NORMAL
- REPL_ACTIVE_DATA is set to TRUE.
- REPL_STATE is set to ENABLED (default)

ALL_DATA_HIGH_PRIORITY (high priority backup, archive or space managed data in filespace)
- REPL_SERVER is set to the default replication server
- REPL_PRIORITY is set to HIGH
- REPL_ACTIVE_DATA is set to FALSE.
- REPL_STATE is set to ENABLED (default)

ACTIVE_DATA_HIGH_PRIORITY (high priority active backup data in filespace only)
- REPL_SERVER is set to the default replication server
- REPL_PRIORITY is set to HIGH
- REPL_ACTIVE_DATA is set to TRUE.
- REPL_STATE is set to ENABLED (default)

Server Default Replication Rules

When the SET REPLSERVER command is issued three server default replication rules are set to the ALL_DATA replication rule. You can update them using one of the following SET commands:

SET BKREPLRuledefault <rule_name>
SET ARREPLRuledefault <rule_name>
SET SPREPLRuledefault <rule_name>

BKREPLRuledefault - specifies the default rule for filespace replication control rules controlling backup data.
ARREPLRuledefault - specifies the default rule for filespace replication control rules controlling archive data.
SPREPLRuledefault - specifies the default rule for filespace replication control rules controlling space managed data.

rule_name

Choices are ALL_DATA, ACTIVE_DATA, ALL_DATA_HIGH_PRIORITY, ACTIVE_DATA_HIGH_PRIORITY or NONE. If NONE is specified and the node and filespace default replication rules are set to DEFAULT, then no data for the filespace is replicated. Essentially its like turning off replication for that data type for all filespace spaces with the rule set to DEFAULT. Rules ACTIVE_DATA and ACTIVE_DATA_HIGH_PRIORITY can only be specified for BKREPLRULEDEFAULT.
*** DEFAULT can not be specified at the server level**

Node Level Default Replication Rules

When a node is enabled for replication, three node level default replication rules are set to the DEFAULT replication rule. You can update them using the DEFINE/UPDATE NODE command:

```
>--+-----------------------------------+--->
   '--BKREPLRuledefault--=--rule_name--'
   '--ARREPLRuledefault--=--rule_name--'
   '--SPREPLRuledefault--=--rule_name--'
```

BKREPLRuledefault - specifies the default rule for filespace replication rule controlling backup data, overrides the server default replication rule setting.
ARREPLRuledefault - specifies the default rule for filespace replication rule controlling archive data, overrides the server default replication rule setting.
SPREPLRuledefault - specifies the default rule for filespace replication rule controlling space managed data, overrides the server default replication rule setting.

rule_name – Choices are ALL_DATA,ACTIVE_DATA,ALL_DATA_HIGH_PRIORITY, ACTIVE_DATA_HIGH_PRIORITY,NONE or DEFAULT. If NONE is specified and the filespace replication rule is set to DEFAULT, then no data for the filespace is replicated. Essentially its like turning off replication for that data type in the filespace. Rules ACTIVE_DATA and ACTIVE_DATA_HIGH_PRIORITY can only be specified for BKREPLRULEDEFAULT. Specifying DEFAULT will cause the filespace replication rule to be set according to the server default replication rules.

Filespace Replication Rules

- When a node is enabled for replication, each data type in every filespace that the node owns is assigned a replication rule of DEFAULT.
  - BKREPLRULE – set to a replication rule to determine replication of backup data in the filespace
  - ARREPLRULE- set to a replication rule to determine replication of archive data in the filespace
  - SPREPLRULE- set to a replication rule to determine replication of space managed data in the filespace

- Each data type in a filespace also has a state which determines if the rule is to be used during replication. Here are the states:
  - ENABLED - this data type is enabled for replication
  - DISABLED - this data type is disabled for replication

- The state can be updated using the UPDATE FILESPACE command and can be set to ENABLED, DISABLED or PURGEDATA. If PURGEDATA is specified, all data for that data type is deleted from the target server and the STATE is set to DISABLED.

Assigning Replication Rules

- Update the server default replication rules with one of the following commands:

  SET BKREPLRuledefault <rule_name>
  SET ARREPLRuledefault <rule_name>
  SET SPREPLRuledefault <rule_name>

- Set the node level default replication rules using the REGISTER/UPDATE NODE command and specify the following parameters:

```
>--+-----------------------------------+---->
   '--BKREPLRuledefault--=--rule_name--'
   '--ARREPLRuledefault--=--rule_name--'
   '--SPREPLRuledefault--=--rule_name--'
```

- Update the filespace replication rules using the UPDATE FILESPACE command.

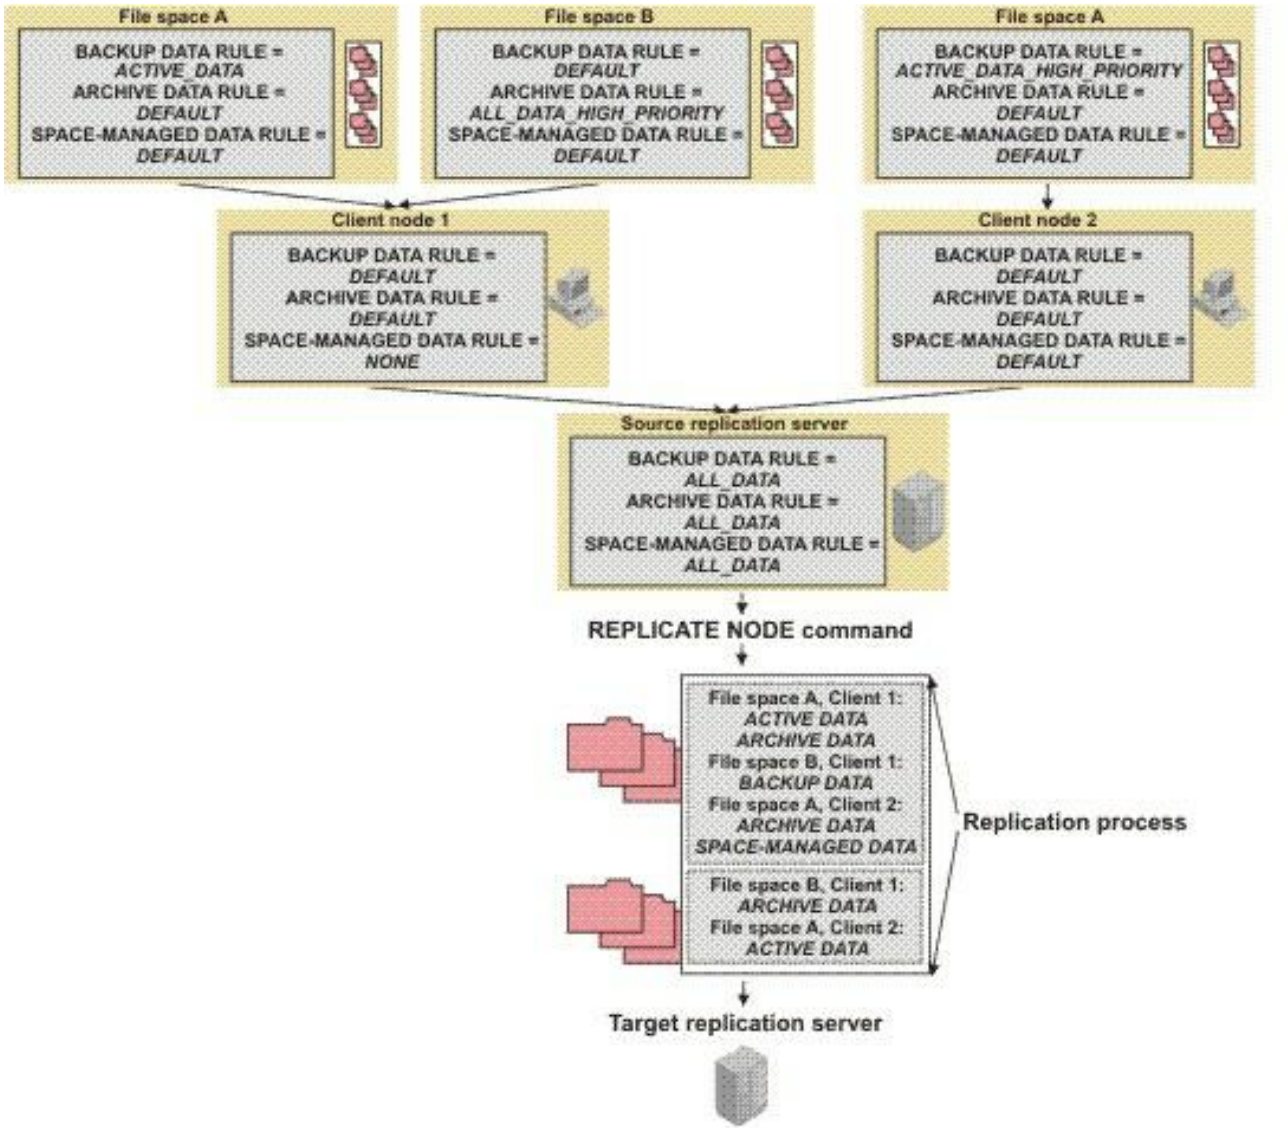

UPDATE FILESPACE COMMAND

```
>>--UPDate FIlespace--node_name--file_space_name-->

      .-NAMEType--=--SERVER------.
   >-+--------------------------+---->
     '-NAMEType--=--+-SERVER--+-'
                    +-UNIcode-+
                    '-FSID----' (1)

      .-CODEType--=--BOTH-----------.
   >--+-----------------------------+-->
      '-CODEType--=--+-UNIcode----+-'
                     +-NONUNIcode-+
                     '-BOTH-------'
>--DATAType--=-------+--BACKup--------+--+-->
               ^     '--ARCHive-------'  |
               ^     '--SPACEManaged--'  |
               ^-----------,------------'
   >--+-----------------------------+->
      '--REPLRule--=--replrule_name--'
   >--+---------------------------------+--><
      '--REPLState--=--+--ENabled----+--'
                       '--DISabled---'
                       '--PURGEdata--'
```

VALIDATE REPLICATION

This command allows the administrator to see at a glance what rules are in effect for filespaces and data types belonging to a node. This is useful to determine if the server, node or filespace replication rules are setup correctly. You can also specify to verify the connection to the target replication server.

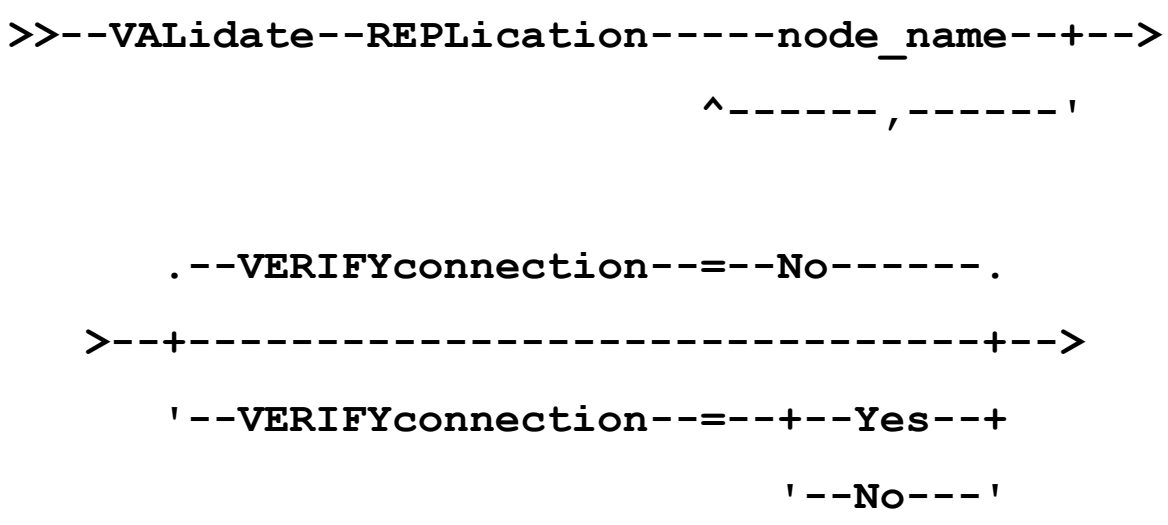

VALIDATE REPLICATION OUTPUT

```
validate replication node1 verifyconnection=yes

                     Node Name: NODE1
                 Filespace Name: \\node1\c$
                           FSID: 1
                           Type: Bkup
Controlling Replication Rule: ACTIVE_DATA
      Replication Rule Level: Server Level
                  Server Name: DRSRV
            Connection Status: Valid Connection

                     Node Name: NODE1
                 Filespace Name: \\node1\c$
                           FSID: 1
                           Type: Arch
Controlling Replication Rule: ALL_DATA_HIGH_PRIORITY
      Replication Rule Level: Node Level
                  Server Name: DRSRV
            Connection Status: Valid Connection
```

Performing Replication

Node Replication

Performing Replication, REPLICATE NODE

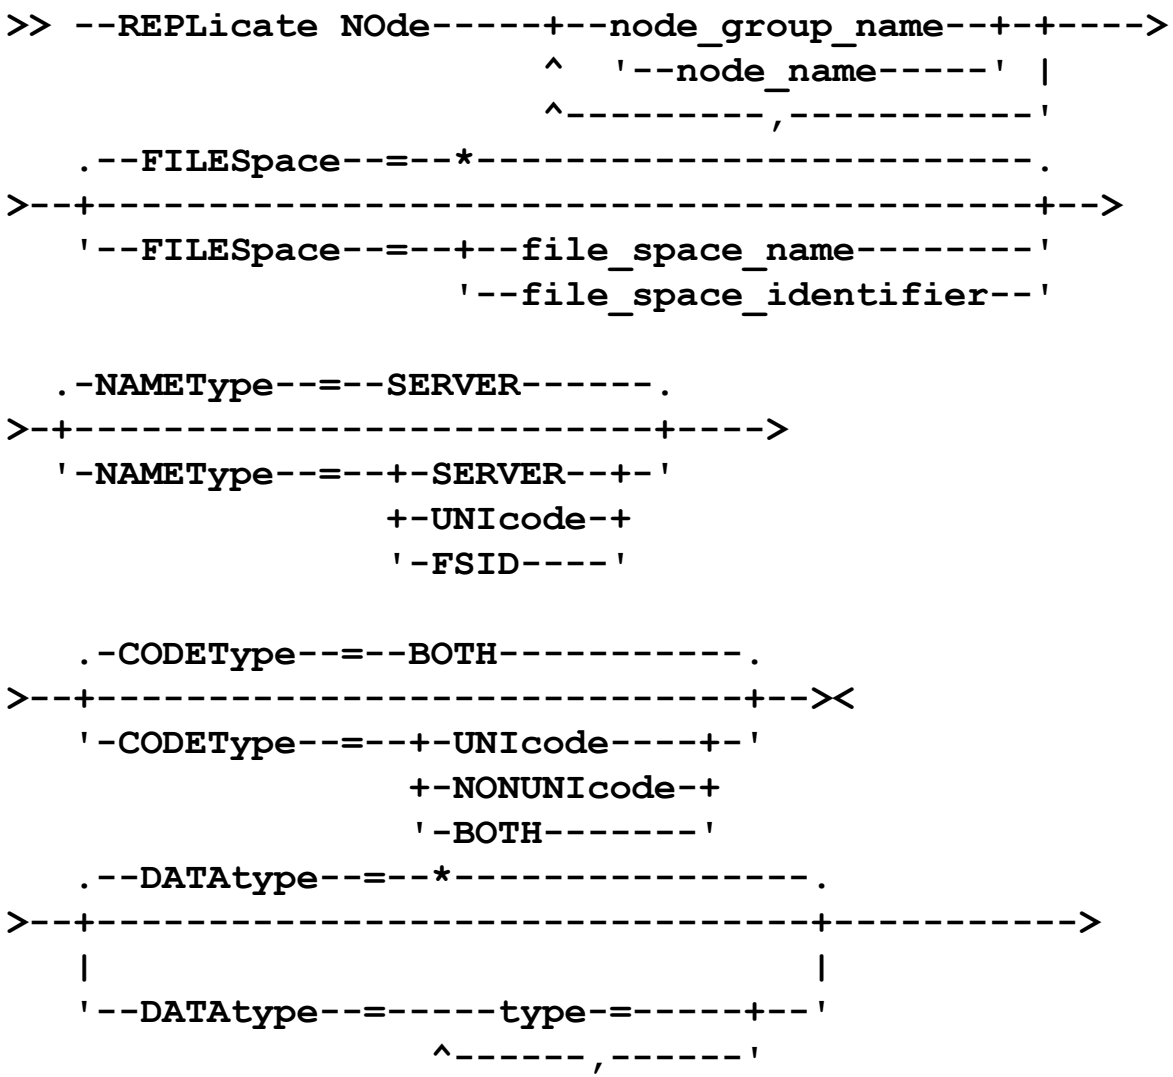

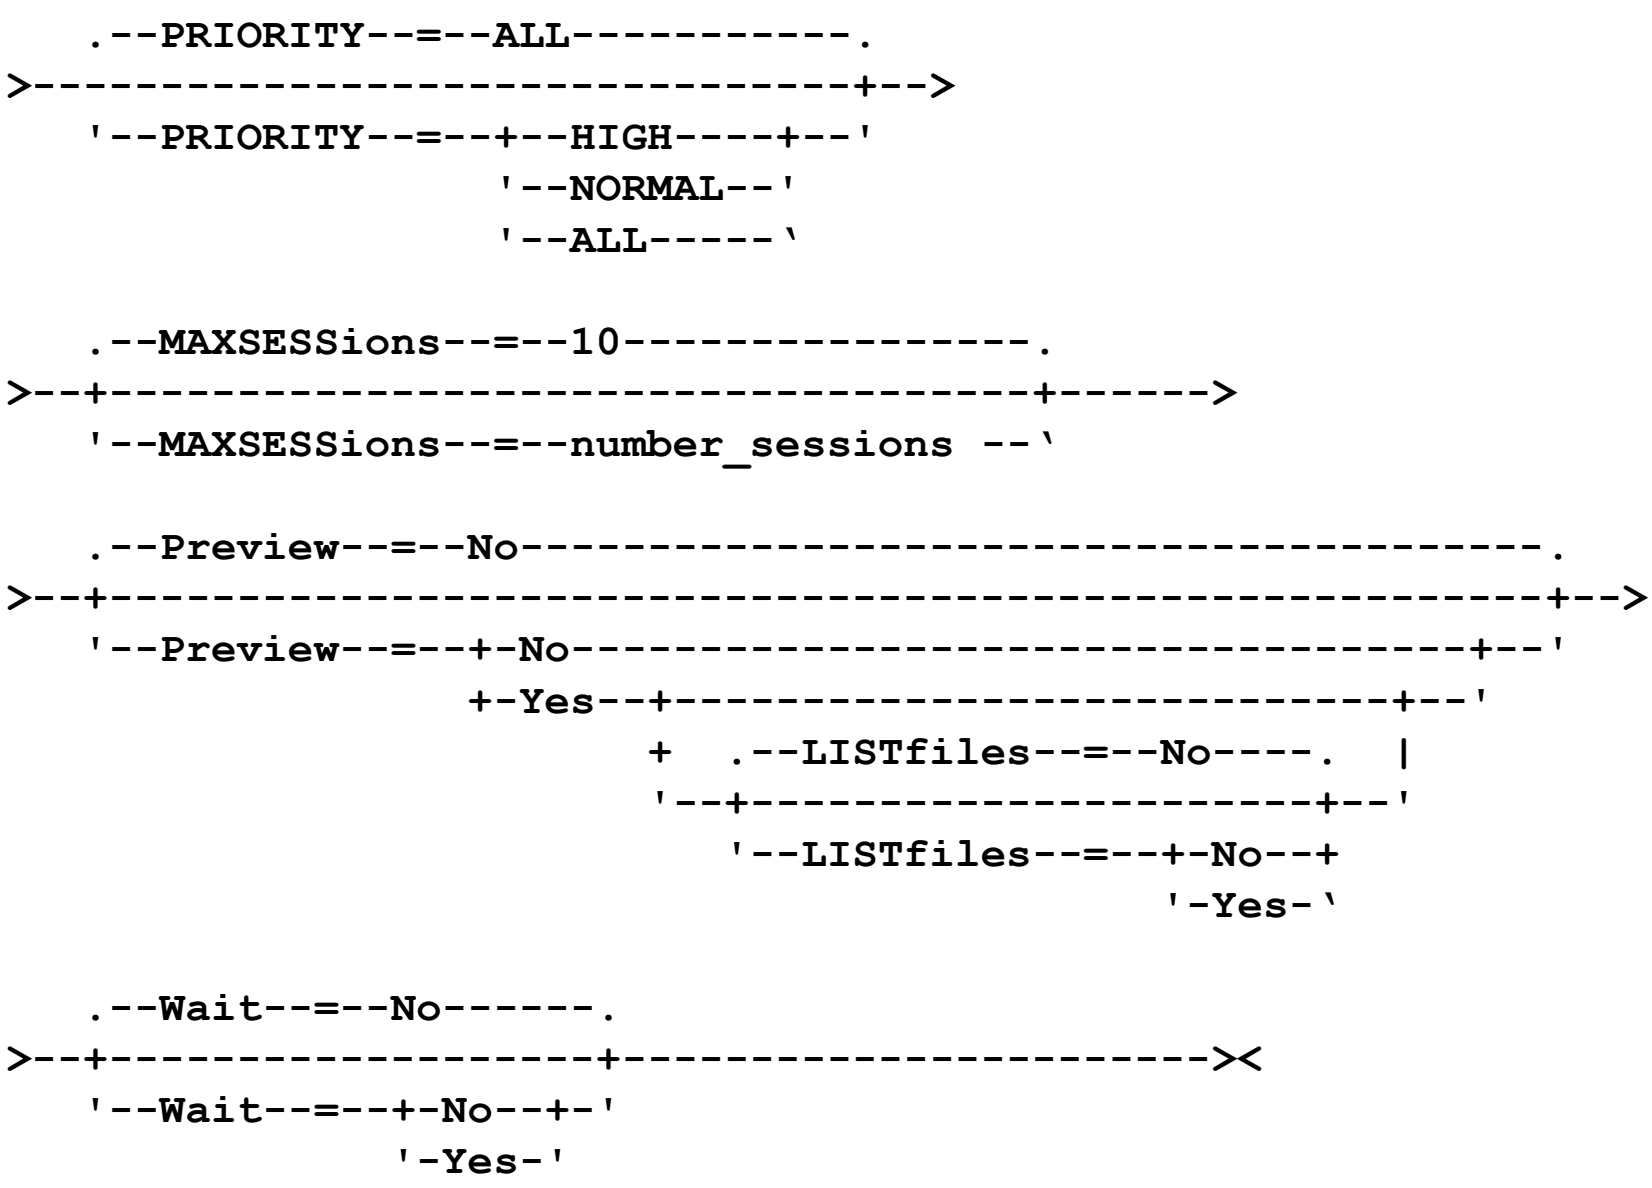

Canceling Replication

Use this command to cancel all node replication processes.

>>-CANcel--REPLication----------><

Disabling/Enabling Sessions (New Parameter)

```
>--SERVer--+-----------------------------------------+-->
           |                .--DIRection--=--Both------.  |
           '--server_name--+--------------------------+--'
                           +--DIRection--=--Both------+
                           +--DIRection--=--INbound---+
                           '--DIRection--=--OUTbound--'
```

New Command: REMove REPLNode

- Removes a node from replication
- Replication configuration and object information in database is deleted for node
- Node is updated to a non-replicating node
- Data is not deleted (source and target)
- Must be issued on all servers that have the node configured for replication

```
>> --REMove REPLNode-----+--node_group_name--+-+---->
                         ^   '--node_name-----' |
                         ^---------,-----------'
```

Q Process Output for Replication (1 of 2)

Process Number	Process Description	Process Status
4	Replicate Node	Replicating Node(s) IRONMAN. File spaces complete: 0. File spaces examining and replicating: 1. File spaces replicating: 0. File spaces not started: 3. Files current: 11,920. Files replicated: 0 of 0. Files updated: 0 of 0. Files deleted: 0 of 0. Bytes Replicated: 0 of 0. Elapsed Time: 0 Day(s), 0 Hour(s), 1 Minute(s).

Q Process Command Output for Replication (cont.)

Replicating Node(s): list of nodes being process by this replication process.

File Spaces complete: number of file spaces that have completed replication.

File spaces examining and replicating: number of file spaces that are determining what files to process and are performing the replication phase.

File spaces replicating: number of file spaces performing only the replication phase.

File spaces not started: number of file spaces that have not yet begun the examining and replicating phases.

Files current: files that do not need replicated, updated or deleted.

Files replicated: the number of files that have been replicated out of a total number to replicate.

Files updated: the number of files that have been updated out of a total number to update.

Files deleted: the number of files that have been deleted out of a total number to delete.

Bytes Replicated: the number of bytes replicated out of a total number of bytes to replicate.

Elapsed Time: the number of days, hours and minutes this process has been running.

Querying Replication

Node Replication

QUERY REPLNODE

Use this command to show the number of client files stored on replication servers.

```
>>--Query REPLNode--node_name--+------------------------------+-->
                               '--target_replication_server--'
```

QUERY REPLNODE OUTPUT

```
tsm: VM-72>q replnode *

Node Name        Type  Filespace Name   FSID    Files on  Replication      Files on
                                            Server  Server (1)       Server (1)
---------------  ----  ---------------  ----  ----------  ---------------  ----------
VM-60            Bkup  /                   1       2,424  VM-71                 2,424
```

- **In the example above the number of files on the server (source) and the number of files on the replication server (target) match which indicates that the node has been replicated and is up to date.**

- **If the replication server name is blank (empty) it indicates that a replication has never been completed. In the example below, the replication server name field and the number of files on the target server will not not show a value or change if the replicate node command was only issued with the preview=yes parameter.**

```
tsm: VM-72>q replnode *

Node Name        Type  Filespace Name   FSID    Files on  Replication      Files on
                                            Server  Server (1)       Server (1)
---------------  ----  ---------------  ----  ----------  ---------------  ----------
VM-60            Bkup  /                   1       2,424
```

QUERYING REPLICATION

```
>>--Query--REPLication--node_name----filespace_name--+-->
                                    ^-------,---------'

                .-NAMEType--=--SERVER------.
             >--+--------------------------+-->
                '-NAMEType--=--+-SERVER--+-'
                               +-UNIcode-+
                               '-FSID----'


              .-CODEType--=--BOTH-----------.
           >--+-----------------------------+-->
              '-CODEType--=--+-UNIcode----+-'
                             +-NONUNIcode-+
                               '-BOTH-------'

             .--DISplay--=--1--------------.
           >--+-----------------------------+-->

              '--DISplay--=--number_of_days--'
```

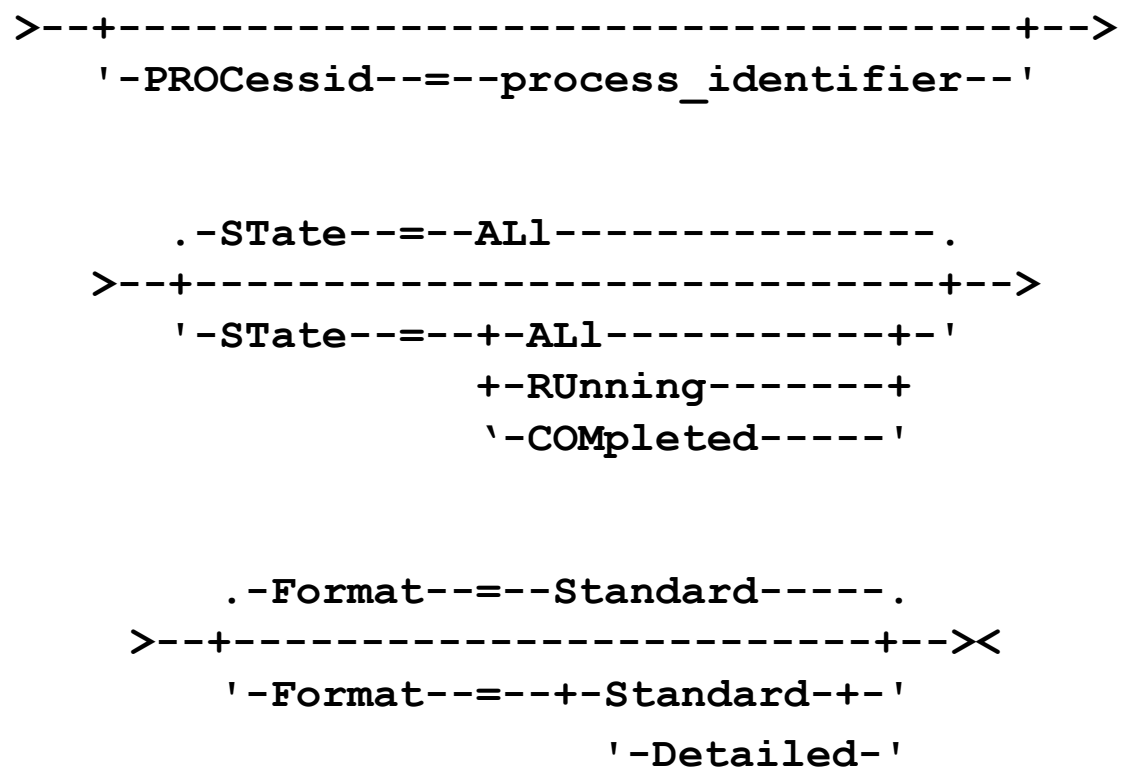

QUERY REPLICATION OUTPUT

query replication payroll 10 nametype=fsid

```
Process  Node     FSID     Start time          State              Phase
-------------------------------------------------------------------------------
5021     payroll  10       07/01/2010          Running  Identifying and replicating
                           15:00:27
4378     payroll  10       07/10/2010          Completed          None
                           15:00:04
```

query replication payroll 10 nametype=fsid f=d

```
                                  Node Name: IRONMAN
                             Filespace Name: /space
                                       FSID: 2
                                 Start Time: 02/08/11   21:44:19
                                   End Time: 02/08/11   21:48:14
                                     Status: Completed
                             Process Number: 4
                                    Command: replicate node ironman
                                      Phase: Completed
                           Process Running Time: 0 Day(s), 0 Hour(s), 4
Minute(s)
                                Completion State: Complete
                        Reason for Incompletion: None
                    Backup Last Update Date/Time:
                             Backup Target Server:
                  Backup Files Needing No Action:           0
                      Backup Files To Replicate:            0
                       Backup Files Replicated:             0
       Backup Files Not Replicated Due To Errors:           0
                 Backup Files Not Yet Replicated:           0
                          Backup Files To Delete:           0
                            Backup Files Deleted:           0
          Backup Files Not Deleted Due To Errors:           0
                          Backup Files To Update:           0
                            Backup Files Updated:           0
          Backup Files Not Updated Due To Errors:           0
```

```
                  Backup Bytes To Replicate (MB):              0
                    Backup Bytes Replicated (MB):              0
                   Backup Bytes Transferred (MB):              0
Backup Bytes Not Replicated Due To Errors (MB):              0
            Backup Bytes Not Yet Replicated (MB):              0
                  Archive Last Update Date/Time: 02/08/11   21:48:14
                          Archive Target Server: NIGLINA
                Archive Files Needing No Action:              0
                     Archive Files To Replicate:         39,416
                       Archive Files Replicated:         39,206
    Archive Files Not Replicated Due To Errors:            210
              Archive Files Not Yet Replicated:              0
                       Archive Files To Delete:              0
                         Archive Files Deleted:              0
       Archive Files Not Deleted Due To Errors:              0
                       Archive Files To Update:              0
                         Archive Files Updated:              0
       Archive Files Not Updated Due To Errors:              0
                Archive Bytes To Replicate (MB):          4,335
                  Archive Bytes Replicated (MB):          4,335
                 Archive Bytes Transferred (MB):              0
Archive Bytes Not Replicated Due To Errors (MB):              0
           Archive Bytes Not Yet Replicated (MB):              0
```

```
          Space Management Last Update Date/Time:
                  Space Management Target Server:
           Space Managed Files Needing No Action:                0
                Space Managed Files To Replicate:                0
                  Space Managed Files Replicated:                0
Space Managed Files Not Replicated Due To Errors:                0
          Space Managed Files Not Yet Replicated:                0
                   Space Managed Files To Delete:                0
                     Space Managed Files Deleted:                0
 Space Managed Files Not Deleted Due To Errors:                  0
                   Space Managed Files To Update:                0
                     Space Managed Files Updated:                0
 Space Managed Files Not Updated Due To Errors:                  0
          Space Managed Bytes To Replicate (MB):                 0
            Space Managed Bytes Replicated (MB):                 0
           Space Managed Bytes Transferred (MB):                 0
Space Managed Bytes Not Replicated Due To Errors (MB):           0
   Space Managed Bytes Not Yet Replicated (MB):                  0
                   Total Files Needing No Action:                0
                        Total Files To Replicate:           39,416
                          Total Files Replicated:           39,206
      Total Files Not Replicated Due To Errors:                210
                Total Files Not Yet Replicated:                  0
                           Total Files To Delete:                0
                             Total Files Deleted:                0
          Total Files Not Deleted Due To Errors:                 0
```

```
                    Total Bytes To Replicate (MB):        4,335
                      Total Bytes Replicated (MB):        4,335
                     Total Bytes Transferred (MB):            0
Total Bytes Not Replicated Due To Errors (MB):                0
              Total Bytes Not Yet Replicated (MB):            0
                    Estimated Percentage Complete:          100
                         Estimated Time Remaining:
                     Estimated Time Of Completion:
```

Setting retention of replication history records

Specifies the retention period for replication records in the server database that will allow you to monitor completed replications. A replication record is created whenever a replication process is started. Replication records for replication processes that have ended are purged during the expiration process. The default is 30 days.

You can adjust the length of time that the server maintains replication information to avoid insufficient or outdated data. The server automatically removes the replication records from the database after the retention period passes.

Syntax

```
>--Set—REPLRETention--number_of_days---->
```

Activity Summary Record for Replication (1 of 2)

```
START_TIME: 2010-07-01 09:27:13.000000
END_TIME: 2010-07-13 09:45:14.00000
ACTIVITY: NODE REPLICATION
NUMBER: 10
ENTITY: PAYROLL
COMMMETH:
 ADDRESS:
SCHEDULE_NAME:
EXAMINED: 239829
AFFECTED: 5348
FAILED: 0
BYTES: 32380928
```

Activity Summary Record for Replication (2 of 2)

IDLE: 0

MEDIAW: 0

PROCESSES: 0

SUCCESSFUL: YES

VOLUME_NAME:

DRIVE_NAME:

LIBRARY_NAME:

LAST_USE:

COMM_WAIT: 0

NUM_OFFSITE_VOLS:

TSM V6.3 Node Replication

Questions?