



Managing Archive Nodes

StorageGRID

NetApp

October 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-115/admin/what-arc-service-is.html> on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Managing Archive Nodes	1
What an Archive Node is	1
What the ARC service is	2
Configuring Archive Node connections to archival storage	2
Archiving to the cloud through the S3 API	2
Archiving to tape through TSM middleware	9
Configuring Archive Node retrieve settings	14
Configuring Archive Node replication	15
Setting Custom alarms for the Archive Node	16
Integrating Tivoli Storage Manager	17
Archive Node configuration and operation	17
Configuration best practices	18
Completing the Archive Node setup	18

Managing Archive Nodes

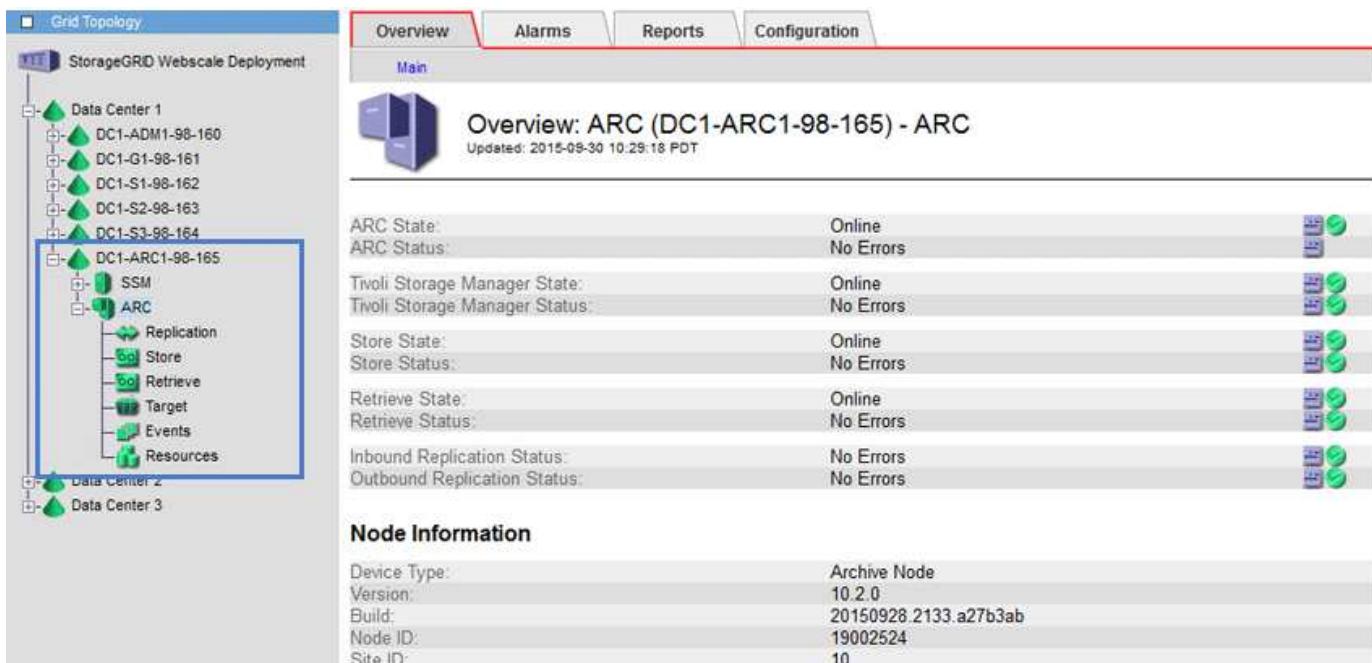
Optionally, each of your StorageGRID system's data center sites can be deployed with an Archive Node, which allows you to connect to a targeted external archival storage system, such as Tivoli Storage Manager (TSM).

After configuring connections to the external target, you can configure the Archive Node to optimize TSM performance, take an Archive Node offline when a TSM server is nearing capacity or unavailable, and configure replication and retrieve settings. You can also set Custom alarms for the Archive Node.

- [What an Archive Node is](#)
- [Configuring Archive Node connections to archival storage](#)
- [Setting Custom alarms for the Archive Node](#)
- [Integrating Tivoli Storage Manager](#)

What an Archive Node is

The Archive Node provides an interface through which you can target an external archival storage system for the long term storage of object data. The Archive Node also monitors this connection and the transfer of object data between the StorageGRID system and the targeted external archival storage system.



The screenshot shows the StorageGRID WebUI interface. The left sidebar displays the 'Grid Topology' with a tree structure of Data Centers (DC1, DC2, DC3) and their sub-components like Storage, Compute, and Network. A specific node under DC1 is selected, highlighted with a blue border. The main content area is titled 'Overview: ARC (DC1-ARC1-98-165) - ARC' and shows the following status information:

ARC State	Online
ARC Status	No Errors
Tivoli Storage Manager State	Online
Tivoli Storage Manager Status	No Errors
Store State	Online
Store Status	No Errors
Retrieve State	Online
Retrieve Status	No Errors
Inbound Replication Status	No Errors
Outbound Replication Status	No Errors

Below this, the 'Node Information' section provides the following details:

Device Type	Archive Node
Version	10.2.0
Build	20150928.2133.a27b3ab
Node ID	19002524
Site ID	10

Object data that cannot be deleted, but is not regularly accessed, can at any time be moved off of a Storage Node's spinning disks and onto external archival storage such as the cloud or tape. This archiving of object data is accomplished through the configuration of a data center site's Archive Node and then the configuration of ILM rules where this Archive Node is selected as the "target" for content placement instructions. The Archive Node does not manage archived object data itself; this is achieved by the external archive device.



Object metadata is not archived, but remains on Storage Nodes.

What the ARC service is

The Archive Node's Archive (ARC) service provides the management interface you can use to configure connections to external archival storage, such as tape through TSM middleware.

It is the ARC service that interacts with an external archival storage system, sending object data for near-line storage and performing retrievals when a client application requests an archived object. When a client application requests an archived object, a Storage Node requests the object data from the ARC service. The ARC service makes a request to the external archival storage system, which retrieves the requested object data and sends it to the ARC service. The ARC service verifies the object data and forwards it to the Storage Node, which in turn returns the object to the requesting client application.

Requests for object data archived to tape through TSM middleware are managed for efficiency of retrievals. Requests can be ordered so that objects stored in sequential order on tape are requested in that same sequential order. Requests are then queued for submission to the storage device. Depending upon the archival device, multiple requests for objects on different volumes can be processed simultaneously.

Configuring Archive Node connections to archival storage

When you configure an Archive Node to connect with an external archive, you must select the target type.

The StorageGRID system supports the archiving of object data to the cloud through an S3 interface or to tape through Tivoli Storage Manager (TSM) middleware.



Once the type of archival target is configured for an Archive Node, the target type cannot be changed.

- [Archiving to the cloud through the S3 API](#)
- [Archiving to tape through TSM middleware](#)
- [Configuring Archive Node retrieve settings](#)
- [Configuring Archive Node replication](#)

Archiving to the cloud through the S3 API

You can configure an Archive Node to connect directly to Amazon Web Services (AWS) or to any other system that can interface to the StorageGRID system through the S3 API.



Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality. The **Cloud Tiering - Simple Storage Service (S3)** option is still supported, but you might prefer to implement Cloud Storage Pools instead.

If you are currently using an Archive Node with the **Cloud Tiering - Simple Storage Service (S3)** option, consider migrating your objects to a Cloud Storage Pool. See the instructions for managing objects with information lifecycle management.

Related information

[Manage objects with ILM](#)

Configuring connection settings for the S3 API

If you are connecting to an Archive Node using the S3 interface, you must configure the connection settings for the S3 API. Until these settings are configured, the ARC service remains in a Major alarm state as it is unable to communicate with the external archival storage system.

 Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality. The **Cloud Tiering - Simple Storage Service (S3)** option is still supported, but you might prefer to implement Cloud Storage Pools instead.

If you are currently using an Archive Node with the **Cloud Tiering - Simple Storage Service (S3)** option, consider migrating your objects to a Cloud Storage Pool. See the instructions for managing objects with information lifecycle management.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have created a bucket on the target archival storage system:
 - The bucket must be dedicated to a single Archive Node. It cannot be used by other Archive Nodes or other applications.
 - The bucket must have the appropriate region selected for your location.
 - The bucket should be configured with versioning suspended.
- Object Segmentation must be enabled and the Maximum Segment Size must be less than or equal to 4.5 GiB (4,831,838,208 bytes). S3 API requests that exceed this value will fail if S3 is used as the external archival storage system.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.



Configuration: ARC (98-127) - Target

Updated: 2015-09-24 15:48:22 PDT

Target Type:

Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:	<input type="text" value="name"/>
Region:	<input type="text" value="Virginia or Pacific Northwest (us-east-1)"/>
Endpoint:	<input type="text" value="https://10.10.10.123:8082"/> <input type="checkbox"/> Use AWS
Endpoint Authentication:	<input type="checkbox"/>
Access Key:	<input type="text" value="ABCD123EFG45AB"/>
Secret Access Key:	<input type="text" value="*****"/>
Storage Class:	<input type="text" value="Standard (Default)"/>

Apply Changes 

4. Select **Cloud Tiering - Simple Storage Service (S3)** from the Target Type drop-down list.



Configuration settings are unavailable until you select a Target Type.

5. Configure the cloud tiering (S3) account through which the Archive Node will connect to the target external S3 capable archival storage system.

Most of the fields on this page are self-explanatory. The following describes fields for which you might need guidance.

- **Region:** Only available if **Use AWS** is selected. The region you select must match the bucket's region.
- **Endpoint and Use AWS:** For Amazon Web Services (AWS), select **Use AWS**. **Endpoint** is then automatically populated with an endpoint URL based on the Bucket Name and Region attributes. For example:

`https://bucket.region.amazonaws.com`

For a non-AWS target, enter the URL of the system hosting the bucket, including the port number. For example:

`https://system.com:1080`

- **End Point Authentication:** Enabled by default. If the network to the external archival storage system is trusted, you can unselect the check box to disable endpoint SSL certificate and hostname verification for the targeted external archival storage system. If another instance of a StorageGRID system is the target archival storage device and the system is configured with publicly signed certificates, you can keep the check box selected.

- **Storage Class:** Select **Standard (Default)** for regular storage. Select **Reduced Redundancy** only for objects that can be easily recreated. **Reduced Redundancy** provides lower cost storage with less reliability. If the targeted archival storage system is another instance of the StorageGRID system, **Storage Class** controls how many interim copies of the object are made at ingest on the target system, if dual commit is used when objects are ingested there.

6. Click **Apply Changes**.

The specified configuration settings are validated and applied to your StorageGRID system. Once configured, the target cannot be changed.

Related information

[Manage objects with ILM](#)

Modifying connection settings for S3 API

After the Archive Node is configured to connect to an external archival storage system through the S3 API, you can modify some settings should the connection change.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

About this task

If you change the Cloud Tiering (S3) account, you must ensure that the user access credentials have read/write access to the bucket, including all objects that were previously ingested by the Archive Node to the bucket.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.



Configuration: ARC (98-127) - Target

Updated: 2015-09-24 15:48:22 PDT

Target Type:

Cloud Tiering - Simple Storage Service (S3)

Cloud Tiering (S3) Account

Bucket Name:	<input type="text" value="name"/>
Region:	<input type="text" value="Virginia or Pacific Northwest (us-east-1)"/>
Endpoint:	<input type="text" value="https://10.10.10.123:8082"/> <input type="checkbox"/> Use AWS
Endpoint Authentication:	<input type="checkbox"/>
Access Key:	<input type="text" value="ABCD123EFG45AB"/>
Secret Access Key:	<input type="text" value="*****"/>
Storage Class:	<input type="text" value="Standard (Default)"/>

Apply Changes



4. Modify account information, as necessary.

If you change the storage class, new object data is stored with the new storage class. Existing object continue to be stored under the storage class set when ingested.



Bucket Name, Region, and Endpoint, use AWS values and cannot be changed.

5. Click **Apply Changes**.

Modifying the Cloud Tiering Service state

You can control the Archive Node's ability read and write to the targeted external archival storage system that connects through the S3 API by changing the state of the Cloud Tiering Service.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The Archive Node must be configured.

About this task

You can effectively take the Archive Node offline by changing the Cloud Tiering Service State to **Read-Write Disabled**.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC**.
3. Select **Configuration > Main**.

Configuration: ARC (98-127) - ARC
Updated: 2015-09-24 17:18:29 PDT

ARC State: Online

Cloud Tiering Service State: Read-Write Enabled

Apply Changes

4. Select a **Cloud Tiering Service State**.
5. Click **Apply Changes**.

Resetting the Store Failure Count for S3 API connection

If your Archive Node connects to an archival storage system through the S3 API, you can reset the Store Failure Count, which can be used to clear the ARVF (Store Failures) alarm.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.

Configuration: ARC (98-127) - Store
Updated: 2015-09-29 17:54:42 PDT

Reset Store Failure Count:

Apply Changes

4. Select **Reset Store Failure Count**.
5. Click **Apply Changes**.

The Store Failures attribute resets to zero.

Migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool

If you are currently using the **Cloud Tiering - Simple Storage Service (S3)** feature to tier object data to an S3 bucket, consider migrating your objects to a Cloud Storage Pool instead. Cloud Storage Pools provide a scalable approach that takes advantage of all of the Storage Nodes in your StorageGRID system.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You have already stored objects in the S3 bucket configured for Cloud Tiering.



Before migrating object data, contact your NetApp account representative to understand and manage any associated costs.

About this task

From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. However, while storage pools consist of Storage Nodes or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external S3 bucket.

Before migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool, you must first create an S3 bucket and then create the Cloud Storage Pool in StorageGRID. Then, you can create a new ILM policy and replace the ILM rule used to store objects in the Cloud Tiering bucket with a cloned ILM rule that stores the same objects in the Cloud Storage Pool.



When objects are stored in a Cloud Storage Pool, copies of those objects cannot also be stored within StorageGRID. If the ILM rule you are currently using for Cloud Tiering is configured to store objects in multiple locations at the same time, consider whether you still want to perform this optional migration because you will lose that functionality. If you continue with this migration, you must create new rules instead of cloning the existing ones.

Steps

1. Create a Cloud Storage Pool.

Use a new S3 bucket for the Cloud Storage Pool to ensure it contains only the data managed by the Cloud Storage Pool.

2. Locate any ILM rules in the active ILM policy that cause objects to be stored in the Cloud Tiering bucket.
3. Clone each of these rules.
4. In the cloned rules, change the placement location to the new Cloud Storage Pool.
5. Save the cloned rules.
6. Create a new policy that uses the new rules.
7. Simulate and activate the new policy.

When the new policy is activated and ILM evaluation occurs, the objects are moved from the S3 bucket configured for Cloud Tiering to the S3 bucket configured for the Cloud Storage Pool. The usable space on

the grid is not affected. After the objects are moved to the Cloud Storage Pool, they are removed from the Cloud Tiering bucket.

Related information

[Manage objects with ILM](#)

Archiving to tape through TSM middleware

You can configure an Archive Node to target a Tivoli Storage Manager (TSM) server that provides a logical interface for storing and retrieving object data to random or sequential access storage devices, including tape libraries.

The Archive Node's ARC service acts as a client to the TSM server, using Tivoli Storage Manager as middleware for communicating with the archival storage system.

TSM management classes

Management classes defined by the TSM middleware outline how the TSM's backup and archive operations function, and can be used to specify rules for content that are applied by the TSM server. Such rules operate independently of the StorageGRID system's ILM policy, and must be consistent with the StorageGRID system's requirement that objects are stored permanently and are always available for retrieval by the Archive Node. After object data is sent to a TSM server by the Archive Node, the TSM lifecycle and retention rules are applied while the object data is stored to tape managed by the TSM server.

The TSM management class is used by the TSM server to apply rules for data location or retention after objects are sent to the TSM server by the Archive Node. For example, objects identified as database backups (temporary content that can be overwritten with newer data) could be treated differently than application data (fixed content that must be retained indefinitely).

Configuring connections to TSM middleware

Before the Archive Node can communicate with Tivoli Storage Manager (TSM) middleware, you must configure a number of settings.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

About this task

Until these settings are configured, the ARC service remains in a Major alarm state as it is unable to communicate with the Tivoli Storage Manager.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.



Configuration: ARC (DC1-ARC1-98-165) - Target

Updated: 2015-09-28 09:56:36 PDT

Target Type:

Tivoli Storage Manager State:

Target (TSM) Account

Server IP or Hostname:	<input type="text" value="10.10.10.123"/>
Server Port:	<input type="text" value="1500"/>
Node Name:	<input type="text" value="ARC-USER"/>
User Name:	<input type="text" value="arc-user"/>
Password:	<input type="text" value="*****"/>
Management Class:	<input type="text" value="sg-mgmtclass"/>
Number of Sessions:	<input type="text" value="2"/>
Maximum Retrieve Sessions:	<input type="text" value="1"/>
Maximum Store Sessions:	<input type="text" value="1"/>

Apply Changes

- From the **Target Type** drop-down list, select **Tivoli Storage Manager (TSM)**.
- For the **Tivoli Storage Manager State**, select **Offline** to prevent retrievals from the TSM middleware server.

By default, the Tivoli Storage Manager State is set to Online, which means that the Archive Node is able to retrieve object data from the TSM middleware server.

- Complete the following information:
 - Server IP or Hostname:** Specify the IP address or fully qualified domain name of the TSM middleware server used by the ARC service. The default IP address is 127.0.0.1.
 - Server Port:** Specify the port number on the TSM middleware server that the ARC service will connect to. The default is 1500.
 - Node Name:** Specify the name of the Archive Node. You must enter the name (arc-user) that you registered on the TSM middleware server.
 - User Name:** Specify the user name the ARC service uses to log in to the TSM server. Enter the default user name (arc-user) or the administrative user you specified for the Archive Node.
 - Password:** Specify the password used by the ARC service to log in to the TSM server.
 - Management Class:** Specify the default management class to use if a management class is not specified when the object is being saved to the StorageGRID system, or the specified management class is not defined on the TSM middleware server.
 - Number of Sessions:** Specify the number of tape drives on the TSM middleware server that are dedicated to the Archive Node. The Archive Node concurrently creates a maximum of one session per mount point plus a small number of additional sessions (less than five).

You must change this value to be the same as the value set for MAXNUMMP (maximum number of mount points) when the Archive Node was registered or updated. (In the register command, the default value of MAXNUMMP used is 1, if no value is set.)

You must also change the value of MAXSESSIONS for the TSM server to a number that is at least as large as the Number of Sessions set for the ARC service. The default value of MAXSESSIONS on the TSM server is 25.

- **Maximum Retrieve Sessions:** Specify the maximum number of sessions that the ARC service can open to the TSM middleware server for retrieve operations. In most cases, the appropriate value is Number of Sessions minus Maximum Store Sessions. If you need to share one tape drive for storage and retrieval, specify a value equal to the Number of Sessions.
- **Maximum Store Sessions:** Specify the maximum number of concurrent sessions that the ARC service can open to the TSM middleware server for archive operations.

This value should be set to one except when the targeted archival storage system is full and only retrievals can be performed. Set this value to zero to use all sessions for retrievals.

7. Click **Apply Changes**.

Optimizing an Archive Node for TSM middleware sessions

You can optimize the performance of an Archive Node that connects to Tivoli Server Manager (TSM) by configuring the Archive Node's sessions.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

About this task

Typically, the number of concurrent sessions that the Archive Node has open to the TSM middleware server is set to the number of tape drives the TSM server has dedicated to the Archive Node. One tape drive is allocated for storage while the rest are allocated for retrieval. However, in situations where a Storage Node is being rebuilt from Archive Node copies or the Archive Node is operating in Read-only mode, you can optimize TSM server performance by setting the maximum number of retrieve sessions to be the same as number of concurrent sessions. The result is that all drives can be used concurrently for retrieval, and, at most, one of these drives can also be used for storage if applicable.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.
4. Change **Maximum Retrieve Sessions** to be the same as **Number of Sessions**.



Configuration: ARC (DC1-ARC1-98-165) - Target

Updated: 2015-09-28 09:56:36 PDT

Target Type:

Tivoli Storage Manager State:

Target (TSM) Account

Server IP or Hostname:	<input type="text" value="10.10.10.123"/>
Server Port:	<input type="text" value="1500"/>
Node Name:	<input type="text" value="ARC-USER"/>
User Name:	<input type="text" value="arc-user"/>
Password:	<input type="text" value="*****"/>
Management Class:	<input type="text" value="sg-mgmtclass"/>
Number of Sessions:	<input type="text" value="2"/>
Maximum Retrieve Sessions:	<input type="text" value="2"/>
Maximum Store Sessions:	<input type="text" value="1"/>

Apply Changes

5. Click **Apply Changes**.

Configuring the archive state and counters for TSM

If your Archive Node connects to a TSM middleware server, you can configure an Archive Node's archive store state to Online or Offline. You can also disable the archive store when the Archive Node first starts up, or reset the failure count being tracked for the associated alarm.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.



Configuration: ARC (DC1-ARC1-98-165) - Store

Updated: 2015-09-29 17:10:12 PDT

Store State

Online

Archive Store Disabled on Startup



Reset Store Failure Count



[Apply Changes](#)



4. Modify the following settings, as necessary:

- Store State: Set the component state to either:
 - Online: The Archive Node is available to process object data for storage to the archival storage system.
 - Offline: The Archive Node is not available to process object data for storage to the archival storage system.
- Archive Store Disabled on Startup: When selected, the Archive Store component remains in the Read-only state when restarted. Used to persistently disable storage to the targeted the archival storage system. Useful when the targeted the archival storage system is unable to accept content.
- Reset Store Failure Count: Reset the counter for store failures. This can be used to clear the ARVF (Stores Failure) alarm.

5. Click **Apply Changes**.

Related information

[Managing an Archive Node when TSM server reaches capacity](#)

Managing an Archive Node when TSM server reaches capacity

The TSM server has no way to notify the Archive Node when either the TSM database or the archival media storage managed by the TSM server is nearing capacity. The Archive Node continues to accept object data for transfer to the TSM server after the TSM server stops accepting new content. This content cannot be written to media managed by the TSM server. An alarm is triggered if this happens. This situation can be avoided through proactive monitoring of the TSM server.

What you'll need

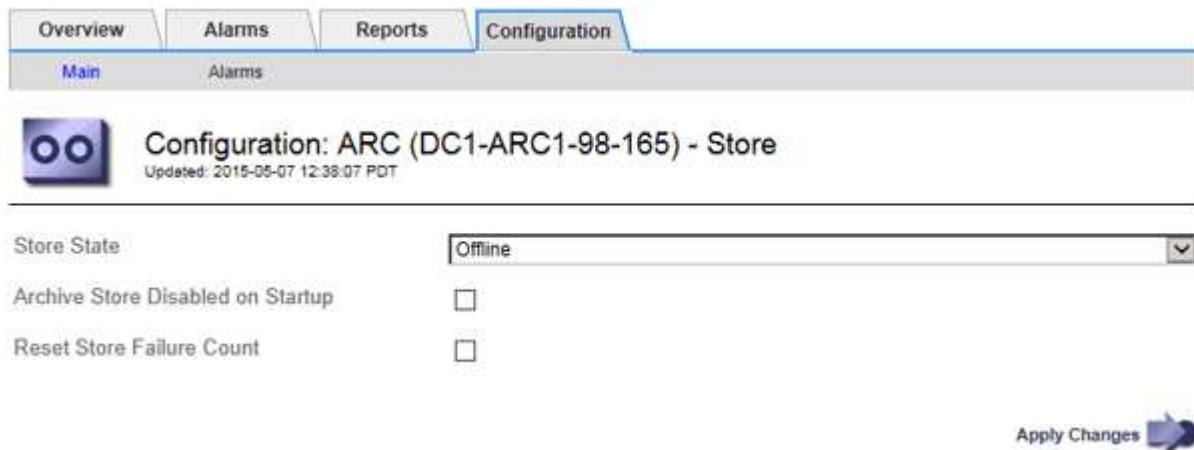
- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

About this task

To prevent the ARC service from sending further content to the TSM server, you can take the Archive Node offline by taking its **ARC > Store** component offline. This procedure can also be useful in preventing alarms when the TSM server is unavailable for maintenance.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.



Configuration: ARC (DC1-ARC1-98-165) - Store
Updated: 2015-05-07 12:39:07 PDT

Store State: Offline

Archive Store Disabled on Startup:

Reset Store Failure Count:

Apply Changes 

4. Change **Store State** to Offline.
5. Select **Archive Store Disabled on Startup**.
6. Click **Apply Changes**.

Setting Archive Node to read-only if TSM middleware reaches capacity

If the targeted TSM middleware server reaches capacity, the Archive Node can be optimized to only perform retrievals.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.
4. Change Maximum Retrieve Sessions to be the same as the number of concurrent sessions listed in Number of Sessions.
5. Change Maximum Store Sessions to 0.



Changing Maximum Store Sessions to 0 is not necessary if the Archive Node is Read-only. Store sessions will not be created.

6. Click **Apply Changes**.

Configuring Archive Node retrieve settings

You can configure the retrieve settings for an Archive Node to set the state to Online or

Offline, or reset the failure counts being tracked for the associated alarms.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Retrieve**.
3. Select **Configuration > Main**.

Configuration: ARC (DC1-ARC1-98-165) - Retrieve
Updated: 2015-05-07 12:24:45 PDT

Retrieve State: Online

Reset Request Failure Count:

Reset Verification Failure Count:

Apply Changes 

4. Modify the following settings, as necessary:

- **Retrieve State:** Set the component state to either:
 - **Online:** The grid node is available to retrieve object data from the archival media device.
 - **Offline:** The grid node is not available to retrieve object data.
- **Reset Request Failures Count:** Select the check box to reset the counter for request failures. This can be used to clear the ARRF (Request Failures) alarm.
- **Reset Verification Failure Count:** Select the check box to reset the counter for verification failures on retrieved object data. This can be used to clear the ARRV (Verification Failures) alarm.

5. Click **Apply Changes**.

Configuring Archive Node replication

You can configure the replication settings for an Archive Node and disable inbound and outbound replication, or reset the failure counts being tracked for the associated alarms.

What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Replication**.

3. Select **Configuration > Main**.

Configuration: ARC (DC1-ARC1-98-165) - Replication
Updated: 2015-05-07 12:21:53 PDT

Reset Inbound Replication Failure Count

Reset Outbound Replication Failure Count

Inbound Replication

Disable Inbound Replication

Outbound Replication

Disable Outbound Replication

Apply Changes

4. Modify the following settings, as necessary:

- **Reset Inbound Replication Failure Count:** Select to reset the counter for inbound replication failures. This can be used to clear the RIRF (Inbound Replications — Failed) alarm.
- **Reset Outbound Replication Failure Count:** Select to reset the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications — Failed) alarm.
- **Disable Inbound Replication:** Select to disable inbound replication as part of a maintenance or testing procedure. Leave cleared during normal operation.

When inbound replication is disabled, object data can be retrieved from the ARC service for replication to other locations in the StorageGRID system, but objects cannot be replicated to this ARC service from other system locations. The ARC service is read-only.

- **Disable Outbound Replication:** Select the check box to disable outbound replication (including content requests for HTTP retrievals) as part of a maintenance or testing procedure. Leave unchecked during normal operation.

When outbound replication is disabled, object data can be copied to this ARC service to satisfy ILM rules, but object data cannot be retrieved from the ARC service to be copied to other locations in the StorageGRID system. The ARC service is write-only.

5. Click **Apply Changes**.

Setting Custom alarms for the Archive Node

You should establish Custom alarms for the ARQL and ARRL attributes that are used to monitor the speed and efficiency of object data retrieval from the archival storage system by the Archive Node.

- ARQL: Average Queue Length. The average time, in microseconds, that object data is queued for retrieval from the archival storage system.

- **ARRL: Average Request Latency.** The average time, in microseconds, needed by the Archive Node to retrieve object data from the archival storage system.

The acceptable values for these attributes depend on how the archival storage system is configured and used. (Go to **ARC > Retrieve > Overview > Main**.) The values set for request timeouts and the number of sessions made available for retrieve requests are particularly influential.

After integration is complete, monitor the Archive Node's object data retrievals to establish values for normal retrieval times and queue lengths. Then, create Custom alarms for ARQL and ARRL that will trigger if an abnormal operating condition arises.

Related information

[Monitor & troubleshoot](#)

Integrating Tivoli Storage Manager

This section includes best practices and set-up information for integrating an Archive Node with a Tivoli Storage Manager (TSM) server, including Archive Node operational details that impact the configuration of the TSM server.

- [Archive Node configuration and operation](#)
- [Configuration best practices](#)
- [Completing the Archive Node setup](#)

Archive Node configuration and operation

Your StorageGRID system manages the Archive Node as a location where objects are stored indefinitely and are always accessible.

When an object is ingested, copies are made to all required locations, including Archive Nodes, based on the Information Lifecycle Management (ILM) rules defined for your StorageGRID system. The Archive Node acts as a client to a TSM server, and the TSM client libraries are installed on the Archive Node by the StorageGRID software installation process. Object data directed to the Archive Node for storage is saved directly to the TSM server as it is received. The Archive Node does not stage object data before saving it to the TSM server, nor does it perform object aggregation. However, the Archive Node can submit multiple copies to the TSM server in a single transaction when data rates warrant.

After the Archive Node saves object data to the TSM server, the object data is managed by the TSM server using its lifecycle/retention policies. These retention policies must be defined to be compatible with the operation of the Archive Node. That is, object data saved by the Archive Node must be stored indefinitely and must always be accessible by the Archive Node, unless it is deleted by the Archive Node.

There is no connection between the StorageGRID system's ILM rules and the TSM server's lifecycle/retention policies. Each operates independently of the other; however, as each object is ingested into the StorageGRID system, you can assign it a TSM management class. This management class is passed to the TSM server along with object data. Assigning different management classes to different object types permits you to configure the TSM server to place object data in different storage pools, or to apply different migration or retention policies as required. For example, objects identified as database backups (temporary content that can be overwritten with newer data) might be treated differently than application data (fixed content that must be retained indefinitely).

The Archive Node can be integrated with a new or an existing TSM server; it does not require a dedicated TSM

server. TSM servers can be shared with other clients, provided that the TSM server is sized appropriately for the maximum expected load. TSM must be installed on a server or virtual machine separate from the Archive Node.

It is possible to configure more than one Archive Node to write to the same TSM server; however, this configuration is only recommended if the Archive Nodes write different sets of data to the TSM server. Configuring more than one Archive Node to write to the same TSM server is not recommended when each Archive Node writes copies of the same object data to the archive. In the latter scenario, both copies are subject to a single point of failure (the TSM server) for what are supposed to be independent, redundant copies of object data.

Archive Nodes do not make use of the Hierarchical Storage Management (HSM) component of TSM.

Configuration best practices

When you are sizing and configuring your TSM server there are best practices you should apply to optimize it to work with the Archive Node.

When sizing and configuring the TSM server, you should consider the following factors:

- Because the Archive Node does not aggregate objects before saving them to the TSM server, the TSM database must be sized to hold references to all objects that will be written to the Archive Node.
- Archive Node software cannot tolerate the latency involved in writing objects directly to tape or other removable media. Therefore, the TSM server must be configured with a disk storage pool for the initial storage of data saved by the Archive Node whenever removable media are used.
- You must configure TSM retention policies to use event-based retention. The Archive Node does not support creation-based TSM retention policies. Use the following recommended settings of `retmin=0` and `retver=0` in the retention policy (which indicates that retention begins when the Archive Node triggers a retention event, and is retained for 0 days after that). However, these values for `retmin` and `retver` are optional.

The disk pool must be configured to migrate data to the tape pool (that is, the tape pool must be the `NXTSTGPOOL` of the disk pool). The tape pool must not be configured as a copy pool of the disk pool with simultaneous write to both pools (that is, the tape pool cannot be a `COPYSTGPOOL` for the disk pool). To create offline copies of the tapes containing Archive Node data, configure the TSM server with a second tape pool that is a copy pool of the tape pool used for Archive Node data.

Completing the Archive Node setup

The Archive Node is not functional after you complete the installation process. Before the StorageGRID system can save objects to the TSM Archive Node, you must complete the installation and configuration of the TSM server and configure the Archive Node to communicate with the TSM server.

For more information about optimizing TSM retrieval and store sessions, see information about managing archival storage.

- [Managing Archive Nodes](#)

Refer to the following IBM documentation, as necessary, as you prepare your TSM server for integration with the Archive Node in a StorageGRID system:

- [IBM Tape Device Drivers Installation and User's Guide](#)
- [IBM Tape Device Drivers Programming Reference](#)

Installing a new TSM server

You can integrate the Archive Node with either a new or an existing TSM server. If you are installing a new TSM server, follow the instructions in your TSM documentation to complete the installation.



An Archive Node cannot be co-hosted with a TSM server.

Configuring the TSM server

This section includes sample instructions for preparing a TSM server following TSM best practices.

The following instructions guide you through the process of:

- Defining a disk storage pool, and a tape storage pool (if required) on the TSM server
- Defining a domain policy that uses the TSM management class for the data saved from the Archive Node, and registering a node to use this domain policy

These instructions are provided for your guidance only; they are not intended to replace TSM documentation, or to provide complete and comprehensive instructions suitable for all configurations. Deployment specific instructions should be provided by a TSM administrator who is familiar both with your detailed requirements, and with the complete set of TSM Server documentation.

Defining TSM tape and disk storage pools

The Archive Node writes to a disk storage pool. To archive content to tape, you must configure the disk storage pool to move content to a tape storage pool.

About this task

For a TSM server, you must define a tape storage pool and a disk storage pool within Tivoli Storage Manager. After the disk pool is defined, create a disk volume and assign it to the disk pool. A tape pool is not required if your TSM server uses disk-only storage.

You must complete a number of steps on your TSM server before you can create a tape storage pool. (Create a tape library and at least one drive in the tape library. Define a path from the server to the library and from the server to the drives, and then define a device class for the drives.) The details of these steps can vary depending upon the hardware configuration and storage requirements of the site. For more information, see the TSM documentation.

The following set of instructions illustrates the process. You should be aware that the requirements for your site could be different depending on the requirements of your deployment. For configuration details and for instructions, see the TSM documentation.



You must log onto the server with administrative privileges and use the `dsmadmc` tool to execute the following commands.

Steps

1. Create a tape library.

```
define library tapelibrary libtype=scsi
```

Where *tapelibrary* is an arbitrary name chosen for the tape library, and the value of libtype can vary depending upon the type of tape library.

2. Define a path from the server to the tape library.

```
define path servername tapelibrary srctype=server desttype=library device=lib-device
```

- *servername* is the name of the TSM server
- *tapelibrary* is the tape library name you defined
- *lib-device* is the device name for the tape library

3. Define a drive for the library.

```
define drive tapelibrary drivename
```

- *drivename* is the name you want to specify for the drive
- *tapelibrary* is the tape library name you defined

You might want to configure an additional drive or drives, depending upon your hardware configuration. (For example, if the TSM server is connected to a Fibre Channel switch that has two inputs from a tape library, you might want to define a drive for each input.)

4. Define a path from the server to the drive you defined.

```
define path servername drivename srctype=server desttype=drive  
library=tapelibrary device=drive-dname
```

- *drive-dname* is the device name for the drive
- *tapelibrary* is the tape library name you defined

Repeat for each drive that you have defined for the tape library, using a separate *drivename* and *drive-dname* for each drive.

5. Define a device class for the drives.

```
define devclass DeviceClassName devtype=lto library=tapelibrary  
format=tapetype
```

- *DeviceClassName* is the name of the device class
- *lto* is the type of drive connected to the server
- *tapelibrary* is the tape library name you defined
- *tapetype* is the tape type; for example, ultrium3

6. Add tape volumes to the inventory for the library.

```
checkin libvolume tapelibrary
```

tapelibrary is the tape library name you defined.

7. Create the primary tape storage pool.

```
define stgpool SGWSTapePool DeviceClassName description=description  
collocate=filespace maxscratch=XX
```

- *SGWSTapePool* is the name of the Archive Node’s tape storage pool. You can select any name for the tape storage pool (as long as the name uses the syntax conventions expected by the TSM server).
- *DeviceClassName* is the name of the device class name for the tape library.
- *description* is a description of the storage pool that can be displayed on the TSM server using the query stgpool command. For example: “Tape storage pool for the Archive Node.”
- *collocate=*filespace** specifies that the TSM server should write objects from the same file space into a single tape.
- XX is one of the following:
 - The number of empty tapes in the tape library (in the case that the Archive Node is the only application using the library).
 - The number of tapes allocated for use by the StorageGRID system (in instances where the tape library is shared).

8. On a TSM server, create a disk storage pool. At the TSM server’s administrative console, enter

```
define stgpool SGWSDiskPool disk description=description  
maxsize=maximum_file_size nextstgpool=SGWSTapePool highmig=percent_high  
lowmig=percent_low
```

- *SGWSDiskPool* is the name of the Archive Node’s disk pool. You can select any name for the disk storage pool (as long as the name uses the syntax conventions expected by the TSM).
- *description* is a description of the storage pool that can be displayed on the TSM server using the query stgpool command. For example, “Disk storage pool for the Archive Node.”
- *maximum_file_size* forces objects larger than this size to be written directly to tape, rather than being cached in the disk pool. It is recommended to set *maximum_file_size* to 10 GB.
- *nextstgpool=*SGWSTapePool** refers the disk storage pool to the tape storage pool defined for the Archive Node.
- *percent_high* sets the value at which the disk pool begins to migrate its contents to the tape pool. It is recommended to set *percent_high* to 0 so that data migration begins immediately
- *percent_low* sets the value at which migration to the tape pool stops. It is recommended to set *percent_low* to 0 to clear out the disk pool.

9. On a TSM server, create a disk volume (or volumes) and assign it to the disk pool.

```
define volume SGWSDiskPool volume_name formatsize=size
```

- *SGWSDiskPool* is the disk pool name.
- *volume_name* is the full path to the location of the volume (for example, /var/local/arc/stage6.dsm) on the TSM server where it writes the contents of the disk pool in

preparation for transfer to tape.

- *size* is the size, in MB, of the disk volume.

For example, to create a single disk volume such that the contents of a disk pool fill a single tape, set the value of *size* to 200000 when the tape volume has a capacity of 200 GB.

However, it might be desirable to create multiple disk volumes of a smaller size, as the TSM server can write to each volume in the disk pool. For example, if the tape size is 250 GB, create 25 disk volumes with a size of 10 GB (10000) each.

The TSM server preallocates space in the directory for the disk volume. This can take some time to complete (more than three hours for a 200 GB disk volume).

Defining a domain policy and registering a node

You need to define a domain policy that uses the TSM management class for the data saved from the Archive Node, and then register a node to use this domain policy.



Archive Node processes can leak memory if the client password for the Archive Node in Tivoli Storage Manager (TSM) expires. Ensure that the TSM server is configured so the client username/password for the Archive Node never expires.

When registering a node on the TSM server for the use of the Archive Node (or updating an existing node), you must specify the number of mount points that the node can use for write operations by specifying the MAXNUMMP parameter to the REGISTER NODE command. The number of mount points is typically equivalent to the number of tape drive heads allocated to the Archive Node. The number specified for MAXNUMMP on the TSM server must be at least as large as the value set for the **ARC > Target > Configuration > Main > Maximum Store Sessions** for the Archive Node, which is set to a value of 0 or 1, as concurrent store sessions are not supported by the Archive Node.

The value of MAXSESSIONS set for the TSM server controls the maximum number of sessions that can be opened to the TSM server by all client applications. The value of MAXSESSIONS specified on the TSM must be at least as large as the value specified for **ARC > Target > Configuration > Main > Number of Sessions** in the Grid Manager for the Archive Node. The Archive Node concurrently creates at most one session per mount point plus a small number (< 5) of additional sessions.

The TSM node assigned to the Archive Node uses a custom domain policy `tsm-domain`. The `tsm-domain` domain policy is a modified version of the “standard” domain policy, configured to write to tape and with the archive destination set to be the StorageGRID system’s storage pool (`SGWSDiskPool`).



You must log in to the TSM server with administrative privileges and use the `dsmadmc` tool to create and activate the domain policy.

Creating and activating the domain policy

You must create a domain policy and then activate it to configure the TSM server to save data sent from the Archive Node.

Steps

1. Create a domain policy.

```
copy domain standard tsm-domain
```

2. If you are not using an existing management class, enter one of the following:

```
define policyset tsm-domain standard  
define mgmtclass tsm-domain standard default
```

default is the default management class for the deployment.

3. Create a copygroup to the appropriate storage pool. Enter (on one line):

```
define copygroup tsm-domain standard default type=archive  
destination=SGWSDiskPool retinit=event retmin=0 retver=0
```

default is the default Management Class for the Archive Node. The values of *retinit*, *retmin*, and *retver* have been chosen to reflect the retention behavior currently used by the Archive Node



Do not set *retinit* to *retinit=create*. Setting *retinit=create* blocks the Archive Node from deleting content since retention events are used to remove content from the TSM server.

4. Assign the management class to be the default.

```
assign defmgmtclass tsm-domain standard default
```

5. Set the new policy set as active.

```
activate policyset tsm-domain standard
```

Ignore the “no backup copy group” warning that appears when you enter the *activate* command.

6. Register a node to use the new policy set on the TSM server. On the TSM server, enter (on one line):

```
register node arc-user arc-password passexp=0 domain=tsm-domain  
MAXNUMMP=number-of-sessions
```

arc-user and *arc-password* are same client node name and password as you define on the Archive Node, and the value of *MAXNUMMP* is set to the number of tape drives reserved for Archive Node store sessions.



By default, registering a node creates an administrative user ID with client owner authority, with the password defined for the node.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.