



# **Archive to tape through TSM middleware**

## **StorageGRID 11.8**

NetApp  
March 19, 2024

# Table of Contents

- Archive to tape through TSM middleware ..... 1
  - TSM management classes ..... 1
  - Configure connections to TSM middleware ..... 1
  - Optimize an Archive Node for TSM middleware sessions ..... 3
  - Configure the archive state and counters for TSM ..... 4
  - Manage an Archive Node when TSM server reaches capacity ..... 5

# Archive 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.

Support for Archive Nodes is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality.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 also deprecated. If you are currently using an Archive Node with this option, [migrate your objects to a Cloud Storage Pool](#) instead.

Additionally, you should remove Archive Nodes from the active ILM policy in StorageGRID 11.7 or earlier. Removing object data stored on Archive Nodes will simplify future upgrades. See [Working with ILM rules and ILM policies](#).

## 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).

## Configure connections to TSM middleware

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

### Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You 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**.

Overview Alarms Reports Configuration

Main Alarms

Configuration: ARC (DC1-ARC1-98-165) - Target  
Updated: 2015-09-28 09:56:36 PDT

Target Type: Tivoli Storage Manager (TSM)

Tivoli Storage Manager State: Online

### Target (TSM) Account

Server IP or Hostname: 10.10.10.123

Server Port: 1500

Node Name: ARC-USER

User Name: arc-user

Password: ●●●●●●

Management Class: sg-mgmtclass

Number of Sessions: 2

Maximum Retrieve Sessions: 1

Maximum Store Sessions: 1

Apply Changes

4. From the **Target Type** drop-down list, select **Tivoli Storage Manager (TSM)**.
5. 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.

6. 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. Select **Apply Changes**.

## Optimize 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.

### Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You 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**.

Overview	Alarms	Reports	Configuration
Main	Alarms		



### 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="password" 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. Select **Apply Changes**.

## Configure 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.

### Before you begin


- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

### Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.

Overview Alarms Reports **Configuration**

Main Alarms


 **Configuration: ARC (DC1-ARC1-98-165) - Store**  
Updated: 2015-09-29 17:10:12 PDT

---

Store State

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 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. Select **Apply Changes**.

#### Related information

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

## Manage 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. This situation can be avoided through proactive monitoring of the TSM server.

#### Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

#### About this task


The Archive Node continues to accept object data for transfer to the TSM server after the TSM server stops accepting new content. This content can't be written to media managed by the TSM server. An alarm is triggered if this happens.

## Prevent ARC service from sending content to TSM server

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**.



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

## Set 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.

### 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. Select **Apply Changes**.



## Copyright information

Copyright © 2024 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—without 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.