



Completing the Archive Node setup

StorageGRID

NetApp
October 03, 2025

Table of Contents

Completing the Archive Node setup	1
Installing a new TSM server	1
Configuring the TSM server	1
Defining TSM tape and disk storage pools	1
Defining a domain policy and registering a node	4
Creating and activating the domain policy	5

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-devicename
```

- *servername* is the name of the TSM server
- *tapelibrary* is the tape library name you defined
- *lib-devicename* 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 SGWTapePool DeviceClassName description=description
collocate=filespace maxscratch=XX
```

- *SGWTapePool* 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=SGWTapePool 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.