



Host-level and middleware procedures

StorageGRID 11.7

NetApp
April 12, 2024

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-117/maintain/linux-migrating-grid-node-to-new-host.html> on April 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Host-level and middleware procedures 1
 - Linux: Migrate grid node to new host 1
 - Archive Node maintenance for TSM middleware 3
 - VMware: Configure virtual machine for automatic restart 8

Host-level and middleware procedures

Linux: Migrate grid node to new host

You can migrate one or more StorageGRID nodes from one Linux host (the *source host*) to another Linux host (the *target host*) to perform host maintenance without impacting the functionality or availability of your grid.

For example, you might want to migrate a node to perform OS patching and reboot.

Before you begin

- You planned your StorageGRID deployment to include migration support.
 - [Node container migration requirements for Red Hat Enterprise Linux or Centos](#)
 - [Node container migration requirements for Ubuntu or Debian](#)
- The target host is already prepared for StorageGRID use.
- Shared storage is used for all per-node storage volumes
- Network interfaces have consistent names across hosts.



In a production deployment, don't run more than one Storage Node on a single host. Using a dedicated host for each Storage Node provides an isolated failure domain.

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same host. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), don't install all instances on the same host.

Export node from source host

As a first step, shut down the grid node and export it from the source Linux host.

Run the following commands on the *source host*.

Steps

1. Obtain the status of all nodes currently running on the source host.

```
sudo storagegrid node status all
```

Example output:

```
Name Config-State Run-State
DC1-ADM1 Configured Running
DC1-ARC1 Configured Running
DC1-GW1 Configured Running
DC1-S1 Configured Running
DC1-S2 Configured Running
DC1-S3 Configured Running
```

2. Identify the name of the node you want to migrate, and stop it if its Run-State is Running.

```
sudo storagegrid node stop DC1-S3
```

Example output:

```
Stopping node DC1-S3
Waiting up to 630 seconds for node shutdown
```

3. Export the node from the source host.

```
sudo storagegrid node export DC1-S3
```

Example output:

```
Finished exporting node DC1-S3 to /dev/mapper/sgws-dc1-s3-var-local.
Use 'storagegrid node import /dev/mapper/sgws-dc1-s3-var-local' if you
want to import it again.
```

4. Make note of the `import` command suggested in the output.

You will run this command on the target host in the next step.

Import node on target host

After exporting the node from the source host, you import and validate the node on the target host. Validation confirms that the node has access to the same block storage and network interface devices as it had on the source host.

Run the following commands on the *target host*.

Steps

1. Import the node on the target host.

```
sudo storagegrid node import /dev/mapper/sgws-dc1-s3-var-local
```

Example output:

```
Finished importing node DC1-S3 from /dev/mapper/sgws-dc1-s3-var-local.
You should run 'storagegrid node validate DC1-S3'
```

2. Validate the node configuration on the new host.

```
sudo storagegrid node validate DC1-S3
```

Example output:

```
Confirming existence of node DC1-S3... PASSED
Checking configuration file /etc/storagegrid/nodes/DC1-S3.conf for node
DC1-S3... PASSED
Checking for duplication of unique values... PASSED
```

3. If any validation errors occur, address them before starting the migrated node.

For troubleshooting information, see the StorageGRID installation instructions for your Linux operating system.

- [Install Red Hat Enterprise Linux or CentOS](#)
- [Install Ubuntu or Debian](#)

Start migrated node

After you validate the migrated node, you start the node by running a command on the *target host*.

Steps

1. Start the node on the new host.

```
sudo storagegrid node start DC1-S3
```

2. Sign in to the Grid Manager and verify that the status of the node is green with no alert.



Verifying that the status of the node is green ensures that the migrated node has fully restarted and rejoined the grid. If the status is not green, don't migrate any additional nodes so that you will not have more than one node out of service.

3. If you are unable to access the Grid Manager, wait for 10 minutes, then run the following command:

```
sudo storagegrid node status _node-name
```

Confirm that the migrated node has a Run-State of Running.

Archive Node maintenance for TSM middleware

Archive Nodes might be configured to target either tape through a TSM middleware server or the cloud through the S3 API. Once configured, an Archive Node's target can't be changed.

If the server hosting the Archive Node fails, replace the server and follow the appropriate recovery procedure.

Fault with archival storage devices

If you determine that there is a fault with the archival storage device that the Archive Node is accessing through Tivoli Storage Manager (TSM), take the Archive Node offline to limit the number of alarms displayed in the StorageGRID system. You can then use the administrative tools of the TSM server or the storage device, or both, to further diagnose and resolve the problem.

Take the Target component offline

Before undertaking any maintenance of the TSM middleware server that might result in it becoming unavailable to the Archive Node, take the Target component offline to limit the number of alarms that are triggered if the TSM middleware server becomes unavailable.

Before you begin

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

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Archive Node > ARC > Target > Configuration > Main**.
3. Change the value of Tivoli Storage Manager State to **Offline**, and click **Apply Changes**.
4. After maintenance is complete, change the value of Tivoli Storage Manager State to **Online**, and click **Apply Changes**.

Tivoli Storage Manager administrative tools

The `dsmadm` tool is the administrative console for the TSM middleware server that is installed on the Archive Node. You can access the tool by typing `dsmadm` at the command line of the server. Log in to the administrative console using the same administrative user name and password that is configured for the ARC service.

The `tsmquery.rb` script was created to generate status information from `dsmadm` in a more readable form. You can run this script by entering the following command at the command line of the Archive Node:

```
/usr/local/arc/tsmquery.rb status
```

For more information about the TSM administrative console `dsmadm`, see the *Tivoli Storage Manager for Linux: Administrator's Reference*.

Object permanently unavailable

When the Archive Node requests an object from the Tivoli Storage Manager (TSM) server and the retrieval fails, the Archive Node retries the request after an interval of 10 seconds. If the object is permanently unavailable (for example, because the object is corrupted on tape), the TSM API has no way to indicate this to the Archive Node, so the Archive Node continues to retry the request.

When this situation occurs, an alarm is triggered, and the value continues to increase. To see the alarm, select **SUPPORT > Tools > Grid topology**. Then, select **Archive Node > ARC > Retrieve > Request Failures**.

If the object is permanently unavailable, you must identify the object and then manually cancel the Archive Node's request as described in the procedure, [Determining if objects are permanently unavailable](#).

A retrieval can also fail if the object is temporarily unavailable. In this case, subsequent retrieval requests should eventually succeed.

If the StorageGRID system is configured to use an ILM rule that creates a single object copy and that copy can't be retrieved, the object is lost and can't be recovered. However, you must still follow the procedure to determine if the object is permanently unavailable to "clean up" the StorageGRID system, to cancel the Archive Node's request, and to purge metadata for the lost object.

Determining if objects are permanently unavailable

You can determine if objects are permanently unavailable by making a request using the TSM administrative console.

Before you begin

- You have specific access permissions.
- You have the `Passwords.txt` file.
- You have the IP address of an Admin Node.

About this task

This example is provided for your information. This procedure can't help you identify all failure conditions that might result in unavailable objects or tape volumes. For information about TSM administration, see TSM Server documentation.

Steps

1. Log in to an Admin Node:
 - a. Enter the following command: `ssh admin@Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
2. Identify the object or objects that could not be retrieved by the Archive Node:
 - a. Go to the directory containing the audit log files: `cd /var/local/audit/export`

The active audit log file is named `audit.log`. Once a day, the active `audit.log` file is saved, and a new `audit.log` file is started. The name of the saved file indicates when it was saved, in the format `yyyy-mm-dd.txt`. After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date.

- b. Search the relevant audit log file for messages indicating that an archived object could not be retrieved. For example, enter: `grep ARCE audit.log | less -n`

When an object can't be retrieved from an Archive Node, the ARCE audit message (Archive Object Retrieve End) displays ARUN (archive middleware unavailable) or GERR (general error) in the result field. The following example line from the audit log shows that the ARCE message terminated with the result ARUN for CBID 498D8A1F681F05B3.

```
[AUDT:[CBID(UI64):0x498D8A1F681F05B3][VLID(UI64):20091127][RSLT(FC32):ARUN][AVER(UI32):7]
[ATIM(UI64):1350613602969243][ATYP(FC32):ARCE][ANID(UI32):13959984][AMID(FC32):ARCI]
[ATID(UI64):4560349751312520631]]
```

For more information see the instructions for understanding audit messages.

- c. Record the CBID of each object that had a request failure.

You might also want to record the following additional information used by the TSM to identify objects saved by the Archive Node:

- **File Space Name:** Equivalent to the Archive Node ID. To find the Archive Node ID, select **SUPPORT > Tools > Grid topology**. Then, select **Archive Node > ARC > Target > Overview**.
- **High Level Name:** Equivalent to the volume ID assigned to the object by the Archive Node. The volume ID takes the form of a date (for example, 20091127), and is recorded as the VLID of the object in archive audit messages.
- **Low Level Name:** Equivalent to the CBID assigned to an object by the StorageGRID system.

d. Log out of the command shell: `exit`

3. Check the TSM server to see if the objects identified in step 2 are permanently unavailable:

a. Log in to the administrative console of the TSM server: `dsmadm`

Use the administrative user name and password that are configured for the ARC service. Enter the user name and password in the Grid Manager. (To see the user name, select **SUPPORT > Tools > Grid topology**. Then, select **Archive Node > ARC > Target > Configuration**.)

b. Determine if the object is permanently unavailable.

For example, you might search the TSM activity log for a data integrity error for that object. The following example shows a search of the activity log for the past day for an object with CBID 498D8A1F681F05B3.

```
> query actlog begindate=-1 search=276C14E94082CC69
12/21/2008 05:39:15 ANR0548W Retrieve or restore
failed for session 9139359 for node DEV-ARC-20 (Bicast ARC)
processing file space /19130020 4 for file /20081002/
498D8A1F681F05B3 stored as Archive - data
integrity error detected. (SESSION: 9139359)
>
```

Depending on the nature of the error, the CBID might not be recorded in the TSM activity log. You might need to search the log for other TSM errors around the time of the request failure.

c. If an entire tape is permanently unavailable, identify the CBIDs for all objects stored on that volume:

```
query content TSM_Volume_Name
```

where `TSM_Volume_Name` is the TSM name for the unavailable tape. The following is an example of the output for this command:

```
> query content TSM-Volume-Name
Node Name      Type Filespace  FSID Client's Name for File Name
-----
DEV-ARC-20     Arch /19130020    216  /20081201/ C1D172940E6C7E12
DEV-ARC-20     Arch /19130020    216  /20081201/ F1D7FBC2B4B0779E
```

The Client's Name for File Name is the same as the Archive Node volume ID (or TSM "high level name") followed by the object's CBID (or TSM "low level name"). That is, the Client's Name for File Name takes the form /Archive Node volume ID /CBID. In the first line of the

example output, the Client's Name for File Name is /20081201/ C1D172940E6C7E12.

Recall also that the Filespace is the node ID of the Archive Node.

You will need the CBID of each object stored on the volume and the node ID of the Archive Node to cancel the retrieval request.

4. For each object that is permanently unavailable, cancel the retrieval request and issue a command to inform the StorageGRID system that the object copy was lost:



Use the ADE Console with caution. If the console is used improperly, it is possible to interrupt system operations and corrupt data. Enter commands carefully, and only use the commands documented in this procedure.

- a. If you aren't already logged in to the Archive Node, log in as follows:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

- b. Access the ADE console of the ARC service: `telnet localhost 1409`

- c. Cancel the request for the object: `/proc/BRTR/cancel -c CBID`

where `CBID` is the identifier of the object that can't be retrieved from the TSM.

If the only copies of the object are on tape, the "bulk retrieval" request is canceled with a message, "1 requests canceled". If copies of the object exist elsewhere in the system, the object retrieval is processed by a different module so the response to the message is "0 requests canceled".

- d. Issue a command to notify the StorageGRID system that an object copy has been lost and that an additional copy must be made: `/proc/CMSI/Object_Lost CBID node_ID`

where `CBID` is the identifier of the object that can't be retrieved from the TSM server, and `node_ID` is the node ID of the Archive Node where the retrieval failed.

You must enter a separate command for each lost object copy: entering a range of CBIDs is not supported.

In most cases, the StorageGRID system immediately begins to make additional copies of object data to ensure that the system's ILM policy is followed.

However, if the ILM rule for the object specified that only one copy be made and that copy has now been lost, the object can't be recovered. In this case running the `Object_Lost` command purges the lost object's metadata from the StorageGRID system.

When the `Object_Lost` command completes successfully, the following message is returned:

```
CLOC_LOST_ANS returned result 'SUCS'
```



The `/proc/CMSI/Object_Lost` command is only valid for lost objects that are stored on Archive Nodes.

- e. Exit the ADE Console: `exit`
 - f. Log out of the Archive Node: `exit`
5. Reset the value of Request Failures in the StorageGRID system:
- a. Go to **Archive Node > ARC > Retrieve > Configuration**, and select **Reset Request Failure Count**.
 - b. Click **Apply Changes**.

Related information

[Administer StorageGRID](#)

[Review audit logs](#)

VMware: Configure virtual machine for automatic restart

If the virtual machine does not restart after VMware vSphere Hypervisor is restarted, you might need to configure the virtual machine for automatic restart.

You should perform this procedure if you notice that a virtual machine does not restart while you are recovering a grid node or performing another maintenance procedure.

Steps

1. In the VMware vSphere Client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Power on**.
3. Configure VMware vSphere Hypervisor to restart the virtual machine automatically in future.

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.