



Configure expanded system

StorageGRID

NetApp
May 08, 2024

Table of Contents

- Configure expanded system. 1
 - Configuration steps after expansion. 1
 - Verify that Storage Node is active 3
 - Copy Admin Node database 3
 - Copy Prometheus metrics 4
 - Copy audit logs 6
 - Rebalance erasure-coded data after adding Storage Nodes. 7

Configure expanded system


Configuration steps after expansion

After completing an expansion, you must perform additional integration and configuration steps.

About this task

You must complete the configuration tasks listed below for the grid nodes you are adding in your expansion. Some tasks might be optional, depending on the options selected when installing and administering your system, and how you want to configure the grid nodes added during the expansion.

Steps

1. If you added a Storage Node, complete the following configuration tasks:
 - a. Review the storage pools used in your ILM rules to ensure the new storage will be used. See [Manage objects with ILM](#).
 - If you added a site, create a storage pool for the site and update ILM rules to use the new storage pool.
 - If you added a Storage Node to an existing site, confirm that the new node uses the correct storage grade.
 - 

By default, a new Storage Node is assigned to the All Storage Nodes storage grade and added to storage pools that use that grade for the site. If you want a new node to use a custom storage grade, you must assign it to the custom grade manually (**ILM > Storage grades**).
 - b. Verify that the Storage Node is ingesting objects. See [Verify that the Storage Node is active](#).
 - c. Rebalance erasure-coded data (only if you were unable to add the recommended number of Storage Nodes). See [Rebalance erasure-coded data after adding Storage Nodes](#).
2. If you added a Gateway Node, complete the following configuration task:
 - If high availability (HA) groups are used for client connections, optionally add the Gateway Node to an HA group. Select **CONFIGURATION > Network > High availability groups** to review the list of existing HA groups and to add the new node. See [Administer StorageGRID](#).
3. If you added an Admin Node, complete the following configuration tasks:
 - a. If single sign-on (SSO) is enabled for your StorageGRID system, create a relying party trust for the new Admin Node. You cannot sign in to the node until you create this relying party trust. See [Configure single sign-on](#).
 - b. If you plan to use the Load Balancer service on Admin Nodes, optionally add the new Admin Node to an HA group. Select **CONFIGURATION > Network > High availability groups** to review the list of existing HA groups and to add the new node. See [Administer StorageGRID](#).
 - c. Optionally, copy the Admin Node database from the primary Admin Node to the expansion Admin Node if you want to keep the attribute and audit information consistent on each Admin Node. See [Copy the Admin Node database](#).
 - d. Optionally, copy the Prometheus database from the primary Admin Node to the expansion Admin Node if you want to keep the historical metrics consistent on each Admin Node. See [Copy Prometheus metrics](#).

- e. Optionally, copy the existing audit logs from the primary Admin Node to the expansion Admin Node if you want to keep the historical log information consistent on each Admin Node. See [Copy audit logs](#).
- f. Optionally, configure access to the system for auditing purposes through an NFS or a CIFS file share. See [Administer StorageGRID](#).



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

- g. Optionally, change the preferred sender for notifications. You can make the expansion Admin Node the preferred sender. Otherwise, an existing Admin Node configured as the preferred sender continues to send notifications, including AutoSupport messages, SNMP notifications, alert emails, and alarm emails (legacy system). See [Administer StorageGRID](#).
4. If you added an Archive Node, complete the following configuration tasks.
 - a. Configure the Archive Node's connection to the targeted external archival storage system. When you complete the expansion, Archive Nodes are in an alarm state until you configure connection information through the **ARC > Target** component. See [Administer StorageGRID](#).
 - b. Update the ILM policy to archive object data through the new Archive Node. See [Manage objects with ILM](#).
 - c. Configure custom alarms for the attributes that are used to monitor the speed and efficiency of object data retrieval from Archive Nodes. See [Administer StorageGRID](#).
 5. To check if expansion nodes were added with an untrusted Client Network or to change whether a node's Client Network is untrusted or trusted, go to **CONFIGURATION > Network > Untrusted Client Network**.

If the Client Network on the expansion node is untrusted, then connections to the node on the Client Network must be made using a load balancer endpoint. See [Administer StorageGRID](#).

6. Configure the Domain Name System (DNS).

If you have been specifying DNS settings separately for each grid node, you must add custom per-node DNS settings for the new nodes. See [Modify DNS configuration for single grid node](#).

The best practice is for the grid-wide DNS server list to contain some DNS servers that are accessible locally from each site. If you just added a new site, add new DNS servers for the site to the grid-wide DNS configuration.



Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see [Modify DNS configuration for single grid node](#).

7. If you added a new site, confirm that Network Time Protocol (NTP) servers are accessible from that site. See [Configure NTP servers](#).



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

Verify that Storage Node is active

After an expansion operation that adds new Storage Nodes completes, the StorageGRID system should automatically start using the new Storage Nodes. You must use the StorageGRID system to verify that the new Storage Node is active.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **NODES > Expansion Storage Node > Storage**.
3. Hover your cursor over the **Storage Used - Object Data** graph to view the value for **Used**, which is the amount of the Total usable space that has been used for object data.
4. Verify that the value of **Used** is increasing as you move your cursor to the right on the graph.

Copy Admin Node database

When adding Admin Nodes through an expansion procedure, you can optionally copy the database from the primary Admin Node to the new Admin Node. Copying the database allows you to retain historical information about attributes, alerts, and alerts.

What you'll need

- You have completed the required expansion steps to add an Admin Node.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

The StorageGRID software activation process creates an empty database for the NMS service on the expansion Admin Node. When the NMS service starts on the expansion Admin Node, it records information for servers and services that are currently part of the system or added later. This Admin Node database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **SUPPORT > Tools > Grid topology** page

To ensure that the Admin Node database is consistent between nodes, you can copy the database from the primary Admin Node to the expansion Admin Node.



Copying the database from the primary Admin Node (*thesource Admin Node*) to an expansion Admin Node can take up to several hours to complete. During this period, the Grid Manager is inaccessible.

Use these steps to stop the MI service and the Management API service on both the primary Admin Node and the expansion Admin Node before copying the database.

Steps

1. Complete the following steps on the primary Admin Node:

- a. Log in to the Admin Node:
 - 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. Run the following command: `recover-access-points`
 - c. Enter the provisioning passphrase.
 - d. Stop the MI service: `service mi stop`
 - e. Stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
2. Complete the following steps on the expansion Admin Node:
- a. Log in to the expansion Admin Node:
 - 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. Stop the MI service: `service mi stop`
 - c. Stop the mgmt-api service: `service mgmt-api stop`
 - d. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - e. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - f. Copy the database from the source Admin Node to the expansion Admin Node:
`/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`
 - g. When prompted, confirm that you want to overwrite the MI database on the expansion Admin Node.

The database and its historical data are copied to the expansion Admin Node. When the copy operation is done, the script starts the expansion Admin Node.
 - h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`
3. Restart the services on the primary Admin Node: `service servermanager start`

Copy Prometheus metrics

After adding a new Admin Node, you can optionally copy the historical metrics maintained by Prometheus from the primary Admin Node to the new Admin Node. Copying the metrics ensures that historical metrics are consistent between Admin Nodes.

What you'll need

- The new Admin Node is installed and running.
- You have the `Passwords.txt` file.

- You have the provisioning passphrase.

About this task

When you add an Admin Node, the software installation process creates a new Prometheus database. You can keep the historical metrics consistent between nodes by copying the Prometheus database from the primary Admin Node (the *source Admin Node*) to the new Admin Node.



Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`
3. Complete the following steps on the new Admin Node:
 - a. Log in to the new Admin Node:
 - 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. Stop the Prometheus service: `service prometheus stop`
 - c. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - d. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - e. Copy the Prometheus database from the source Admin Node to the new Admin Node:
`/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
 - f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the new Admin Node.

The original Prometheus database and its historical data are copied to the new Admin Node. When the copy operation is done, the script starts the new Admin Node. The following status appears:

```
Database cloned, starting services
```

- g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:

```
ssh-add -D
```

4. Restart the Prometheus service on the source Admin Node.

```
service prometheus start
```

Copy audit logs

When you add a new Admin Node through an expansion procedure, its AMS service only logs events and actions that occur after it joins the system. As required, you can copy audit logs from a previously installed Admin Node to the new expansion Admin Node so that it is in sync with the rest of the StorageGRID system.

What you'll need

- You have completed the required expansion steps to add an Admin Node.
- You have the `Passwords.txt` file.

About this task

To make historical audit messages available on a new Admin Node, you must copy the audit log files manually from an existing Admin Node to the expansion Admin Node.



By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:

- You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
- You explicitly specified that audit messages should be saved only on the local nodes that generated them.

See [Configure audit messages and log destinations](#) for details.

Steps

1. Log in to the primary Admin Node:

- a. Enter the following command: `ssh admin@_primary_Admin_Node_IP`
- b. Enter the password listed in the `Passwords.txt` file.
- c. Enter the following command to switch to root: `su -`
- d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

2. Stop the AMS service to prevent it from creating a new file: `service ams stop`

3. Rename the `audit.log` file to ensure that it does not overwrite the file on the expansion Admin Node you are copying it to:

```
cd /var/local/audit/export
ls -l
mv audit.log new_name.txt
```

4. Copy all audit log files to the expansion Admin Node:

```
scp -p * IP_address:/var/local/audit/export
```

5. If prompted for the passphrase for `/root/.ssh/id_rsa`, enter the SSH Access Password for the Primary

Admin Node listed in the `Passwords.txt` file.

6. Restore the original `audit.log` file:

```
mv new_name.txt audit.log
```

7. Start the AMS service:

```
service ams start
```

8. Log out from the server:

```
exit
```

9. Log in to the expansion Admin Node:

- a. Enter the following command: `ssh admin@expansion_Admin_Node_IP`
- b. Enter the password listed in the `Passwords.txt` file.
- c. Enter the following command to switch to root: `su -`
- d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

10. Update the user and group settings for the audit log files:

```
cd /var/local/audit/export
chown ams-user:bycast *
```

11. Log out from the server:

```
exit
```

Rebalance erasure-coded data after adding Storage Nodes

In some cases, you might need to rebalance erasure-coded data after you add new Storage Nodes.

What you'll need

- You have completed the expansion steps to add the new Storage Nodes.
- You have reviewed the [considerations for rebalancing erasure-coded data](#).



Only perform this procedure if the **Low Object Storage** alert has been triggered for one or more Storage Nodes at a site and you were unable to add the recommended number of new Storage Nodes.

- You understand that replicated object data will not be moved by this procedure and that the EC rebalance procedure does not consider the replicated data usage on each Storage Node when determining where to move erasure-coded data.
- You have the `Passwords.txt` file.

About this task

When the EC rebalance procedure is running, the performance of ILM operations and S3 and Swift client operations are likely to be impacted. For this reason, you should only perform this procedure in limited cases.



The EC rebalance procedure temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the rebalance is complete. If there is not enough storage for the reservation, the EC rebalance procedure will fail. Storage reservations are released when the EC rebalance procedure completes, whether the procedure failed or succeeded.



S3 and Swift API operations to upload objects (or object parts) might fail during the EC rebalancing procedure if they require more than 24 hours to complete. Long-duration PUT operations will fail if the applicable ILM rule uses Strict or Balanced placement on ingest. The following error will be reported:

```
500 Internal Server Error
```

Steps

1. Review the current object storage details for the site you plan to rebalance.
 - a. Select **NODES**.
 - b. Select the first Storage Node at the site.
 - c. Select the **Storage** tab.
 - d. Hover your cursor over the Storage Used - Object Data chart to see the current amount of replicated data and erasure-coded data on the Storage Node.
 - e. Repeat these steps to view the other Storage Nodes at the site.
2. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

3. Enter the following command:

```
rebalance-data start --site "site-name"
```

For `"site-name"`, specify the first site where you added new Storage Node or nodes. Enclose `site-name` in quotes.

The EC rebalance procedure starts, and a job ID is returned.

4. Copy the job ID.
5. Monitor the status of the EC rebalance procedure.
 - To view the status of a single EC rebalance procedure:

```
rebalance-data status --job-id job-id
```

For *job-id*, specify the ID that was returned when you started the procedure.

- To view the status of the current EC rebalance procedure and any previously completed procedures:

```
rebalance-data status
```



To get help on the rebalance-data command:

```
rebalance-data --help
```

- To view the estimated time to completion and the completion percentage for the current job, select **SUPPORT > Tools > Metrics**. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

6. Perform additional steps, based on the status returned:

- If the status is `In progress`, the EC rebalance operation is still running. You should periodically monitor the procedure until it completes.
- If the status is `Failure`, perform the [failure steps](#).
- If the status is `Success`, perform the [success step](#).

7. If the EC rebalance procedure is generating too much load (for example, ingest operations are affected), pause the procedure.

```
rebalance-data pause --job-id job-id
```

8. If you need to terminate the EC rebalance procedure (for example, so you can perform a StorageGRID software upgrade), enter the following:

```
rebalance-data terminate --job-id job-id
```



When you terminate an EC rebalance procedure, any data fragments that have already been moved remain in the new location. Data is not moved back to the original location.

9. If the status of the EC rebalance procedure is `Failure`, follow these steps:

- a. Confirm that all Storage Nodes at the site are connected to the grid.
- b. Check for and resolve any alerts that might be affecting these Storage Nodes.

For information about specific alerts, see the monitoring and troubleshooting instructions.

- c. Restart the EC rebalance procedure:

```
rebalance-data start --job-id job-id
```

- d. If the status of the EC rebalance procedure is still `Failure`, contact technical support.

10. If the status of the EC rebalance procedure is `Success`, optionally [review object storage](#) to see the updated details for the site.

Erasure-coded data should now be more balanced among the Storage Nodes at the site.

11. If you are using erasure coding at more than one site, run this procedure for all other affected sites.

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.