Restoring and validating grid nodes

StorageGRID

NetApp
October 03, 2025

This PDF was generated from https://docs.netapp.com/us-en/storagegrid-115/maintain/fixing-mssing-network-interface-errors.html on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Restoring and validating grid nodes 1
Fixing missing network interface errors 2
Fixing missing block device errors 3

Restoring and validating grid nodes

You must restore the grid configuration files for any failed grid nodes, and then validate the grid configuration files and resolve any errors.

About this task

You can import any grid node that should be present on the host, as long as its `/var/local` volume was not lost as a result of the failure of the previous host. For example, the `/var/local` volume might still exist if you used shared storage for StorageGRID system data volumes, as described in the StorageGRID installation instructions for your Linux operating system. Importing the node restores its node configuration file to the host.

If it is not possible to import missing nodes, you must recreate their grid configuration files.

You must then validate the grid configuration file, and resolve any networking or storage issues that might occur before going on to restart StorageGRID. When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering.

See the installation instructions for more information on the location of the `/var/local` volume for a node.

Steps

1. At the command line of the recovered host, list all currently configured StorageGRID grid nodes:`sudo storagegrid node list`

   If no grid nodes are configured, there will be no output. If some grid nodes are configured, expect output in the following format:

   ```
   Name               Metadata-Volume
   ===================================================================
   dc1-adm1           /dev/mapper/sgws-adm1-var-local
   dc1-gw1            /dev/mapper/sgws-gw1-var-local
   dc1-sn1            /dev/mapper/sgws-sn1-var-local
   dc1-arc1           /dev/mapper/sgws-arc1-var-local
   ```

   If some or all of the grid nodes that should be configured on the host are not listed, you need to restore the missing grid nodes.

2. To import grid nodes that have a `/var/local` volume:

   a. Run the following command for each node you want to import:`sudo storagegrid node import node-var-local-volume-path`

      The `storagegrid node import` command succeeds only if the target node was shut down cleanly on the host on which it last ran. If that is not the case, you will observe an error similar to the following:

      `This node (node-name) appears to be owned by another host (UUID host-uuid).`

      `Use the --force flag if you are sure import is safe.`

   b. If you see the error about the node being owned by another host, run the command again with the `--force` flag to complete the import:`sudo storagegrid --force node import node-var-`

`local-volume-path`

> Any nodes imported with the `--force` flag will require additional recovery steps before they can rejoin the grid, as described in “Performing additional recovery steps, if required.”

3. For grid nodes that do not have a `/var/local` volume, recreate the node’s configuration file to restore it to the host.

   Follow the guidelines in “Creating node configuration files” in the installation instructions.

   > When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering. For Linux deployments, ensure that the configuration file name contains the node name. You should use the same network interfaces, block device mappings, and IP addresses when possible. This practice minimizes the amount of data that needs to be copied to the node during recovery, which could make the recovery significantly faster (in some cases, minutes rather than weeks).

   > If you use any new block devices (devices that the StorageGRID node did not use previously) as values for any of the configuration variables that start with `BLOCK_DEVICE_` when you are recreating the configuration file for a node, be sure to follow all of the guidelines in “Fixing missing block device errors.”

4. Run the following command on the recovered host to list all StorageGRID nodes.

   ```
   sudo storagegrid node list
   ```

5. Validate the node configuration file for each grid node whose name was shown in the storagegrid node list output:

   ```
   sudo storagegrid node validate node-name
   ```

   You must address any errors or warnings before starting the StorageGRID host service. The following sections give more detail on errors that might have special significance during recovery.

Related information

Install Red Hat Enterprise Linux or CentOS

Install Ubuntu or Debian

Fixing missing network interface errors

Fixing missing block device errors

What’s next: Performing additional recovery steps, if required

Fixing missing network interface errors

If the host network is not configured correctly or a name is misspelled, an error occurs when StorageGRID checks the mapping specified in the `/etc/storagegrid/nodes/node-name.conf` file.

You might see an error or warning matching this pattern:

`Checking configuration file `/etc/storagegrid/nodes/node-name.conf` for node *node-name*...` `ERROR: node-name: GRID_NETWORK_TARGET = host-interface-name` *node-name*: Interface '*host-interface-name*' does not exist`

The error could be reported for the Grid Network, the Admin Network, or the Client Network. This error means that the `/etc/storagegrid/nodes/node-name.conf` file maps the indicated StorageGRID network to the host interface named `host-interface-name`, but there is no interface with that name on the current host.

If you receive this error, verify that you completed the steps in “Deploying new Linux hosts.” Use the same names for all host interfaces as were used on the original host.

If you are unable to name the host interfaces to match the node configuration file, you can edit the node configuration file and change the value of the GRID_NETWORK_TARGET, the ADMIN_NETWORK_TARGET, or the CLIENT_NETWORK_TARGET to match an existing host interface.

Make sure the host interface provides access to the appropriate physical network port or VLAN, and that the interface does not directly reference a bond or bridge device. You must either configure a VLAN (or other virtual interface) on top of the bond device on the host, or use a bridge and virtual Ethernet (veth) pair.

Related information

Deploying new Linux hosts

Fixing missing block device errors

The system checks that each recovered node maps to a valid block device special file or a valid softlink to a block device special file. If StorageGRID finds invalid mapping in the `/etc/storagegrid/nodes/node-name.conf` file, a missing block device error displays.

If you observe an error matching this pattern:

`Checking configuration file /etc/storagegrid/nodes/node-name.conf for node node-name... ERROR: node-name: BLOCK_DEVICE_PURPOSE = path-name` *node-name*: *path-name* does not exist`

It means that `/etc/storagegrid/nodes/node-name.conf` maps the block device used by *node-name* for PURPOSE to the given path-name in the Linux file system, but there is not a valid block device special file, or softlink to a block device special file, at that location.

Verify that you completed the steps in “Deploying new Linux hosts.” Use the same persistent device names for all block devices as were used on the original host.

If you are unable to restore or recreate the missing block device special file, you can allocate a new block device of the appropriate size and storage category and edit the node configuration file to change the value of BLOCK_DEVICE_PURPOSE to point to the new block device special file.

Determine the appropriate size and storage category from the tables in the “Storage requirements” section of the installation instructions for your Linux operating system. Review the recommendations in “Configuring host storage” before proceeding with the block device replacement.

If you must provide a new block storage device for any of the configuration file variables starting with `BLOCK_DEVICE_` because the original block device was lost with the failed host, ensure the new block device is unformatted before attempting further recovery procedures. The new block device will be unformatted if you are using shared storage and have created a new volume. If you are unsure, run the following command against any new block storage device special files.

Run the following command only for new block storage devices. Do not run this command if you believe the block storage still contains valid data for the node being recovered, as any data on the device will be lost.

```
sudo dd if=/dev/zero of=/dev/mapper/my-block-device-name bs=1G count=1
```

Related information

Deploying new Linux hosts

Install Red Hat Enterprise Linux or CentOS

Install Ubuntu or Debian

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—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.