



Adding storage volumes to Storage Nodes

StorageGRID 11.5

NetApp
August 30, 2024

Table of Contents

- Adding storage volumes to Storage Nodes 1
 - VMware: Adding storage volumes to a Storage Node 2
 - Linux: Adding direct-attached or SAN volumes to a Storage Node 4

Adding storage volumes to Storage Nodes

You can expand the storage capacity of Storage Nodes that have 16 or fewer storage volumes by adding additional storage volumes. You might need to add storage volumes to more than one Storage Node to satisfy ILM requirements for replicated or erasure-coded copies.

What you'll need

Before adding storage volumes, review the guidelines for adding storage capacity to ensure that you know where to add volumes to meet the requirements of your ILM policy.

Adding storage capacity



These instructions apply to software-based Storage Nodes only. See the installation and maintenance instructions for the SG6060 appliance to learn how to add storage volumes to SG6060 by installing expansion shelves. Other appliance Storage Nodes cannot be expanded.

[SG6000 storage appliances](#)

About this task

The underlying storage of a Storage Node is divided into a number of storage volumes. Storage volumes are block-based storage devices that are formatted by the StorageGRID system and mounted to store objects. Each Storage Node can support up to 16 storage volumes, which are called *object stores* in the Grid Manager.



Object metadata is always stored in object store 0.

Each object store is mounted on a volume that corresponds to its ID. That is, the object store with an ID of 0000 corresponds to the `/var/local/rangedb/0` mount point.

Before adding new storage volumes, use the Grid Manager to view the current object stores for each Storage Node as well as the corresponding mount points. You can use this information when adding storage volumes.

Steps

1. Select **Nodes** > *site* > **Storage Node** > **Storage**.
2. Scroll down to view the amounts of available storage for each volume and object store.

For appliance Storage Nodes, the Worldwide Name for each disk matches the volume world-wide identifier (WWID) that appears when you view standard volume properties in SANtricity software (the management software connected to the appliance's storage controller).

To help you interpret disk read and write statistics related to volume mount points, the first portion of the name shown in the **Name** column of the Disk Devices table (that is, *sd**c*, *sd**d*, *sd**e*, and so on) matches the value shown in the **Device** column of the Volumes table.

Disk Devices						
Name	World Wide Name	I/O Load		Read Rate		Write Rate
croot(8:1,sda1)	N/A	0.03%		0 bytes/s		4 KB/s
cvloc(8:2,sda2)	N/A	0.37%		0 bytes/s		29 KB/s
sdc(8:16,sdb)	N/A	0.00%		0 bytes/s		0 bytes/s
sdd(8:32,sdc)	N/A	0.00%		0 bytes/s		183 bytes/s
sde(8:48,sdd)	N/A	0.00%		0 bytes/s		12 bytes/s

Volumes						
Mount Point	Device	Status	Size	Available		Write Cache Status
/	croot	Online	10.50 GB	3.46 GB		Unknown
/var/local	cvloc	Online	96.59 GB	94.99 GB		Unknown
/var/local/rangedb/0	sdc	Online	53.66 GB	53.57 GB		Enabled
/var/local/rangedb/1	sdd	Online	53.66 GB	53.57 GB		Enabled
/var/local/rangedb/2	sde	Online	53.66 GB	53.57 GB		Enabled

Object Stores						
ID	Size	Available		Object Data	Object Data (%)	Health
0000	53.66 GB	48.21 GB		976.25 KB		0.00%
0001	53.66 GB	53.57 GB		0 bytes		0.00%
0002	53.66 GB	53.57 GB		0 bytes		0.00%

3. Follow the instructions for your platform to add new storage volumes to the Storage Node.

- [VMware: Adding storage volumes to a Storage Node](#)
- [Linux: Adding direct-attached or SAN volumes to a Storage Node](#)

VMware: Adding storage volumes to a Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by using VMware vSphere to add volumes.

What you'll need

- You must have access to the instructions for installing StorageGRID for VMware deployments.
- You must have the `Passwords.txt` file.
- You must have specific access permissions.



Do not attempt to add storage volumes to a Storage Node while a software upgrade, recovery procedure, or another expansion procedure is active.

About this task

The Storage Node is unavailable for a brief time when you add storage volumes. You should perform this procedure on one Storage Node at a time to avoid impacting client-facing grid services.

Steps

1. If necessary, install new storage hardware and create new VMware datastores.
2. Add one or more hard disks to the virtual machine for use as storage (object stores).
 - a. Open VMware vSphere Client.
 - b. Edit the virtual machine settings to add one or more additional hard disks.

The hard disks are typically configured as Virtual Machine Disks (VMDKs). VMDKs are more commonly used and are easier to manage, while RDMS may provide better performance for workloads that use larger object sizes (for example, greater than 100 MB). For more information about adding hard disks to virtual machines, see the VMware vSphere documentation.

3. Restart the virtual machine by using the **Restart Guest OS** option in the VMware vSphere Client, or by entering the following command in an ssh session to the virtual machine:`sudo reboot`



Do not use **Power Off** or **Reset** to restart the virtual machine.

4. Configure the new storage for use by the Storage Node:
 - a. Log in to the grid 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. When you are logged in as root, the prompt changes from `$` to `#`.

- b. Configure the new storage volumes:

```
sudo add_rangedbs.rb
```

This script finds any new storage volumes and prompts you to format them.

- c. Enter **y** to accept the formatting.
- d. If any of the volumes have previously been formatted, decide if you want to reformat them.
 - Enter **y** to reformat.
 - Enter **n** to skip reformatting. The storage volumes are formatted.
- e. When asked, enter **y** to stop storage services.

The storage services are stopped, and the `setup_rangedbs.sh` script runs automatically. After the volumes are ready for use as rangedbs, the services start again.

5. Check that the services start correctly:
 - a. View a listing of the status of all services on the server:

```
sudo storagegrid-status
```

The status is updated automatically.

- b. Wait until all services are Running or Verified.
- c. Exit the status screen:

Ctrl+C

6. Verify that the Storage Node is online:
 - a. Sign in to the Grid Manager using a supported browser.
 - b. Select **Support > Tools > Grid Topology**.
 - c. Select **site > Storage Node > LDR > Storage**.
 - d. Select the **Configuration** tab and then the **Main** tab.
 - e. If the **Storage State - Desired** drop-down list is set to Read-only or Offline, select **Online**.
 - f. Click **Apply Changes**.
7. To see the new object stores:
 - a. Select **Nodes > site > Storage Node > Storage**.
 - b. View the details in the **Object Stores** table.

Result

You can now use the expanded capacity of the Storage Nodes to save object data.

Related information

[Install VMware](#)

Linux: Adding direct-attached or SAN volumes to a Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by adding new block storage devices, making them visible to the Linux hosts, and adding the new block device mappings to the StorageGRID configuration file used for the Storage Node.

What you'll need

- You must have access to the instructions for installing StorageGRID for your Linux platform.
- You must have the `Passwords.txt` file.
- You must have specific access permissions.



Do not attempt to add storage volumes to a Storage Node while a software upgrade, recovery procedure, or another expansion procedure is active.

About this task

The Storage Node is unavailable for a brief time when you add storage volumes. You should perform this procedure on one Storage Node at a time to avoid impacting client-facing grid services.

Steps

1. Install the new storage hardware.

For more information, see the documentation provided by your hardware vendor.

2. Create new block storage volumes of the desired sizes.
 - Attach the new disk drives and update the RAID controller configuration as needed, or allocate the new SAN LUNs on the shared storage arrays and allow the Linux host to access them.
 - Use the same persistent naming scheme you used for the storage volumes on the existing Storage Node.
 - If you use the StorageGRID node migration feature, make the new volumes visible to other Linux hosts that are migration targets for this Storage Node. For more information, see the instructions for installing StorageGRID for your Linux platform.
3. Log into the Linux host supporting the Storage Node as root or with an account that has sudo permission.
4. Confirm that the new storage volumes are visible on the Linux host.

You might have to rescan for devices.

5. Run the following command to temporarily disable the Storage Node:

```
sudo storagegrid node stop <node-name>
```

6. Using a text editor such as vim or pico, edit the node configuration file for the Storage Node, which can be found at `/etc/storagegrid/nodes/<node-name>.conf`.
7. Locate the section of the node configuration file that contains the existing object storage block device mappings.

In the example, `BLOCK_DEVICE_RANGEDB_00` to `BLOCK_DEVICE_RANGEDB_03` are the existing object storage block device mappings.

```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

8. Add new object storage block device mappings corresponding to the block storage volumes you added for this Storage Node.

Make sure to start at the next `BLOCK_DEVICE_RANGEDB_nn`. Do not leave a gap.

- Based on the example above, start at BLOCK_DEVICE_RANGEDB_04.
- In the example below, four new block storage volumes have been added to the node:
BLOCK_DEVICE_RANGEDB_04 to BLOCK_DEVICE_RANGEDB_07.

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-snl-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-snl-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-snl-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-snl-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-snl-rangedb-3
<strong>BLOCK_DEVICE_RANGEDB_04 = /dev/mapper/sgws-snl-rangedb-4</strong>
<strong>BLOCK_DEVICE_RANGEDB_05 = /dev/mapper/sgws-snl-rangedb-5</strong>
<strong>BLOCK_DEVICE_RANGEDB_06 = /dev/mapper/sgws-snl-rangedb-6</strong>
<strong>BLOCK_DEVICE_RANGEDB_07 = /dev/mapper/sgws-snl-rangedb-7</strong>
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

```

9. Run the following command to validate your changes to the node configuration file for the Storage Node:

```
sudo storagegrid node validate <node-name>
```

Address any errors or warnings before proceeding to the next step.

If you observe an error similar to the following, it means that the node configuration file is attempting to map 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.



```

Checking configuration file for node <node-name>...
ERROR: BLOCK_DEVICE_<PURPOSE> = <path-name>
<path-name> is not a valid block device

```

Verify that you entered the correct <path-name>.

10. Run the following command to restart the node with the new block device mappings in place:

```
sudo storagegrid node start <node-name>
```


11. Log in to the Storage Node as admin using the password listed in the `Passwords.txt` file.

12. Check that the services start correctly:

a. View a listing of the status of all services on the server:

```
sudo storagegrid-status
```

The status is updated automatically.

b. Wait until all services are Running or Verified.

c. Exit the status screen:

```
Ctrl+C
```

13. Configure the new storage for use by the Storage Node:

a. Configure the new storage volumes:

```
sudo add_rangedbs.rb
```

This script finds any new storage volumes and prompts you to format them.

b. Enter **y** to format the storage volumes.

c. If any of the volumes have previously been formatted, decide if you want to reformat them.

- Enter **y** to reformat.

- Enter **n** to skip reformatting. The storage volumes are formatted.

d. When asked, enter **y** to stop storage services.

The storage services are stopped, and the `setup_rangedbs.sh` script runs automatically. After the volumes are ready for use as rangedbs, the services start again.

14. Check that the services start correctly:

a. View a listing of the status of all services on the server:

```
sudo storagegrid-status
```

The status is updated automatically.

b. Wait until all services are Running or Verified.

c. Exit the status screen:

```
Ctrl+C
```

15. Verify that the Storage Node is online:

a. Sign in to the Grid Manager using a supported browser.

b. Select **Support > Tools > Grid Topology**.

c. Select **site > Storage Node > LDR > Storage**.

d. Select the **Configuration** tab and then the **Main** tab.

e. If the **Storage State - Desired** drop-down list is set to Read-only or Offline, select **Online**.

f. Click **Apply Changes**.

16. To see the new object stores:

- a. Select **Nodes > site > Storage Node > Storage**.
- b. View the details in the **Object Stores** table.

Result

You can now use the expanded capacity of the Storage Nodes to save object data.

Related information

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

[Install Ubuntu or Debian](#)

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.