



Manage disks

ONTAP 9

NetApp
April 24, 2024

Table of Contents

- Manage disks 1
 - Overview of managing disks 1
 - How hot spare disks work 1
 - How low spare warnings can help you manage your spare disks 2
 - Additional root-data partitioning management options 2
 - When you need to update the Disk Qualification Package 3
 - Disk and partition ownership 3
 - Remove a failed disk 21
 - Disk sanitization 21
 - Commands for managing disks 28
 - Commands for displaying space usage information 29
 - Commands for displaying information about storage shelves 30

Manage disks

Overview of managing disks

You can perform various procedures to manage disks in your system.

- **Aspects of managing disks**
 - [When you need to update the Disk Qualification Package](#)
 - [How hot spare disks work](#)
 - [How low spare warnings can help you manage your spare disks](#)
 - [Additional root-data partitioning management options](#)
- **Disk and partition ownership**
 - [Disk and partition ownership](#)
- **Failed disk removal**
 - [Remove a failed disk](#)
- **Disk sanitization**
 - [Disk sanitization](#)

How hot spare disks work

A hot spare disk is a disk that is assigned to a storage system and is ready for use, but is not in use by a RAID group and does not hold any data.

If a disk failure occurs within a RAID group, the hot spare disk is automatically assigned to the RAID group to replace the failed disks. The data of the failed disk is reconstructed on the hot spare replacement disk in the background from the RAID parity disk. The reconstruction activity is logged in the `/etc/message` file and an AutoSupport message is sent.

If the available hot spare disk is not the same size as the failed disk, a disk of the next larger size is chosen and then downsized to match the size of the disk that it is replacing.

Spare requirements for multi-disk carrier disk

Maintaining the proper number of spares for disks in multi-disk carriers is critical for optimizing storage redundancy and minimizing the amount of time that ONTAP must spend copying disks to achieve an optimal disk layout.

You must maintain a minimum of two hot spares for multi-disk carrier disks at all times. To support the use of the Maintenance Center and to avoid issues caused by multiple concurrent disk failures, you should maintain at least four hot spares for steady state operation, and replace failed disks promptly.

If two disks fail at the same time with only two available hot spares, ONTAP might not be able to swap the contents of both the failed disk and its carrier mate to the spare disks. This scenario is called a stalemate. If this happens, you are notified through EMS messages and AutoSupport messages. When the replacement carriers become available, you must follow the instructions that are provided by the EMS messages. For more information, see Knowledge Base article [RAID Layout Cannot Be Autocorrected - AutoSupport message](#)

How low spare warnings can help you manage your spare disks

By default, warnings are issued to the console and logs if you have fewer than one hot spare drive that matches the attributes of each drive in your storage system.

You can change the threshold value for these warning messages to ensure that your system adheres to best practices.

About this task

You should set the “min_spare_count” RAID option to “2” to ensure that you always have the minimum recommended number of spare disks.

Step

1. Set the option to “2”:

```
storage raid-options modify -node nodename -name min_spare_count -value 2
```

Additional root-data partitioning management options

Beginning with ONTAP 9.2, a new root-data partitioning option is available from the Boot Menu that provides additional management features for disks that are configured for root-data partitioning.

The following management features are available under the Boot Menu Option 9.

- **Unpartition all disks and remove their ownership information**

This option is useful if your system is configured for root-data partitioning and you need to reinitialize it with a different configuration.

- **Clean configuration and initialize node with partitioned disks**

This option is useful for the following:

- Your system is not configured for root-data partitioning and you would like to configure it for root-data partitioning
- Your system is incorrectly configured for root-data partitioning and you need to correct it
- You have an AFF platform or a FAS platform with only SSDs attached that is configured for the previous version of root-data partitioning and you want to upgrade it to the newer version of root-data partitioning to gain increased storage efficiency

- **Clean configuration and initialize node with whole disks**

This option is useful if you need to:

- Unpartition existing partitions
- Remove local disk ownership
- Reinitialize your system with whole disks using RAID-DP

When you need to update the Disk Qualification Package

The Disk Qualification Package (DQP) adds full support for newly qualified drives. Before you update drive firmware or add new drive types or sizes to a cluster, you must update the DQP. A best practice is to also update the DQP regularly; for example, every quarter or semi-annually.

You need to download and install the DQP in the following situations:

- Whenever you add a new drive type or size to the node

For example, if you already have 1-TB drives and add 2-TB drives, you need to check for the latest DQP update.

- Whenever you update the disk firmware
- Whenever newer disk firmware or DQP files are available
- Whenever you upgrade to a new version of ONTAP.

The DQP is not updated as part of an ONTAP upgrade.

Related information

[NetApp Downloads: Disk Qualification Package](#)

[NetApp Downloads: Disk Drive Firmware](#)

Disk and partition ownership

Disk and partition ownership

You can manage the ownership of disks and partitions.

You can perform the following tasks:

- [Display disk and partition ownership](#)

You can view disk ownership to determine which node controls the storage. You can also view the partition ownership on systems that use shared disks.

- [Change settings for automatic assignment of disk ownership](#)

You can select a non-default policy for automatically assigning disk ownership or disable automatic assignment of disk ownership.

- [Manually assign ownership of unpartitioned disks](#)

If your cluster is not configured to use automatic disk ownership assignment, you must assign ownership manually.

- [Manually assign ownership of partitioned disks](#)

You can set the ownership of the container disk or the partitions manually or by using auto-assignment—

just as you do for unpartitioned disks.

- **Remove a failed disk**

A disk that has failed completely is no longer considered by ONTAP to be a usable disk, and you can immediately disconnect the disk from the shelf.

- **Remove ownership from a disk**

ONTAP writes disk ownership information to the disk. Before you remove a spare disk or its shelf from a node, you should remove its ownership information so that it can be properly integrated into another node.

About automatic assignment of disk ownership

The automatic assignment of unowned disks is enabled by default. Automatic disk ownership assignments occur 10 minutes after HA pair initialization and every five minutes during normal system operation.

When you add a new disk to an HA pair, for example, when replacing a failed disk, responding to a “low spares” message, or adding capacity, the default auto-assignment policy assigns ownership of the disk to a node as a spare.

The default auto-assignment policy is based on platform-specific characteristics, or the DS460C shelf if your HA pair has only these shelves, and it uses one of the following methods (policies) to assign disk ownership:

Assignment method	Effect on node assignments	Platform configurations that default to the assignment method
bay	Even-numbered bays are assigned to node A and odd-numbered bays to node B.	Entry-level systems in an HA pair configuration with a single, shared shelf.
shelf	All disks in the shelf are assigned to node A.	Entry-level systems in an HA pair configuration with one stack of two or more shelves, and MetroCluster configurations with one stack per node, two or more shelves.
split shelf This policy falls under the “default” value for the <code>-autoassign -policy</code> parameter of the <code>storage disk option</code> command for applicable platform and shelf configurations.	Disks on the left side of the shelf are assigned to node A and on the right side to Node B. Partial shelves on HA pairs are shipped from the factory with disks populated from the shelf edge toward the center.	Most AFF platforms and some MetroCluster configurations.
stack	All disks in the stack are assigned to node A.	Stand-alone entry-level systems and all other configurations.

<p>half-drawer</p> <p>This policy falls under the “default” value for the <code>-autoassign-policy</code> parameter of the <code>storage disk option</code> command for applicable platform and shelf configurations.</p>	<p>All drives in the left half of a DS460C drawer (drive bays 0 to 5) are assigned to node A; all drives in the right half of a drawer (drive bays 6 to 11) are assigned to node B.</p> <p>When initializing an HA pair with only DS460C shelves, automatic assignment of disk ownership is not supported. You must manually assign ownership for drives containing root/container drives that have the root partition by conforming to the half-drawer policy.</p>	<p>HA pairs with only DS460C shelves, after HA pair initialization (boot up).</p> <p>After an HA pair boots up, automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the root drives/container drives that have the root partition) and any drives added in the future.</p> <p>If your HA pair has DS460C shelves in addition to other shelf models, the half-drawer policy is not used. The default policy used is dictated by platform-specific characteristics.</p>
---	---	---

Auto-assignment settings and modifications:

- You can display the current auto-assignment settings (on/off) with the `storage disk option show` command.
- You can disable automatic assignment by using the `storage disk option modify` command.
- If the default auto-assignment policy is not desirable in your environment, you can specify (change) the bay, shelf, or stack assignment method using the `-autoassign-policy` parameter in the `storage disk option modify` command.

Learn how to [Change settings for automatic assignment of disk ownership](#).



The half-drawer and split-shelf default auto-assignment policies are unique because they cannot be set by users like the bay, shelf, and stack policies can.

In Advanced Drive Partitioning (ADP) systems, to make auto-assign work on half-populated shelves, drives must be installed in the correct shelf bays based on what type of shelf you have:

- If your shelf is not a DS460C shelf, install drives equally on the far left side and far right side moving toward the middle. For example, six drives in bays 0-5 and six drives in bays 18-23 of a DS224C shelf.
- If your shelf is a DS460C shelf, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer. For the remaining drives, evenly distribute them across each drawer by filling drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.



If drives are not installed in the correct shelf bays on half-populated shelves, when a container drive fails and is replaced, ONTAP does not auto-assign ownership. In this case, assignment of the new container drive needs to be done manually. After you have assigned ownership for the container drive, ONTAP automatically handles any drive partitioning and partitioning assignments that are required.

In some situations where auto-assignment will not work, you need to manually assign disk ownership using the `storage disk assign` command:

- If you disable auto-assignment, new disks are not available as spares until they are manually assigned to a node.
- If you want disks to be auto-assigned and you have multiple stacks or shelves that must have different ownership, one disk must have been manually assigned on each stack or shelf so that automatic ownership assignment works on each stack or shelf.
- If auto-assignment is enabled and you manually assign a single drive to a node that isn't specified in the active policy, auto-assignment stops working and an EMS message is displayed.

Learn how to [Manually assign disk ownership of unpartitioned disks](#).

Learn how to [Manually assign disk ownership of partitioned disks](#).

Display disk and partition ownership

You can view disk ownership to determine which node controls the storage. You can also view the partition ownership on systems that use shared disks.

Steps

1. Display the ownership of physical disks:

```
storage disk show -ownership
```

```
cluster::> storage disk show -ownership
Disk      Aggregate Home      Owner    DR Home  Home ID      Owner ID  DR
Home ID  Reserver  Pool
-----  -
-----  -
1.0.0    aggr0_2  node2    node2    -      2014941509  2014941509  -
2014941509 Pool0
1.0.1    aggr0_2  node2    node2    -      2014941509  2014941509  -
2014941509 Pool0
1.0.2    aggr0_1  node1    node1    -      2014941219  2014941219  -
2014941219 Pool0
1.0.3    -         node1    node1    -      2014941219  2014941219  -
2014941219 Pool0
```

2. If you have a system that uses shared disks, you can display the partition ownership:

```
storage disk show -partition-ownership
```

```

cluster::> storage disk show -partition-ownership

```

Container	Container	Root	Data
Disk	Aggregate	Root Owner	Data Owner
Owner ID		Owner ID	Owner ID
1.0.0	-	node1	node1
1886742616		1886742616	1886742616
1.0.1	-	node1	node1
1886742616		1886742616	1886742616
1.0.2	-	node2	node2
1886742657		1886742657	1886742657
1.0.3	-	node2	node2
1886742657		1886742657	1886742657

Change settings for automatic assignment of disk ownership

You can use the `storage disk option modify` command to select a non-default policy for automatically assigning disk ownership or to disable automatic assignment of disk ownership.

Learn about [automatic assignment of disk ownership](#).

About this task

If you have an HA pair with only DS460C shelves, the default auto-assignment policy is half-drawer. You cannot change to a non-default policy (bay, shelf, stack).

Steps

1. Modify automatic disk assignment:

- a. If you want to select a non-default policy, enter:

```
storage disk option modify -autoassign-policy autoassign_policy -node node_name
```

- Use `stack` as the `autoassign_policy` to configure automatic ownership at the stack or loop level.
- Use `shelf` as the `autoassign_policy` to configure automatic ownership at the shelf level.
- Use `bay` as the `autoassign_policy` to configure automatic ownership at the bay level.

- b. If you want to disable automatic disk ownership assignment, enter:

```
storage disk option modify -autoassign off -node node_name
```

2. Verify the automatic assignment settings for the disks:

```
storage disk option show
```

```
cluster1::> storage disk option show
```

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
cluster1-1	on	on	on	default
cluster1-2	on	on	on	default

Manually assign disk ownership of unpartitioned disks

If your HA pair is not configured to use automatic disk ownership assignment, you must manually assign ownership. If you are initializing an HA pair that has only DS460C shelves, you must manually assign ownership for the root drives.

About this task

- If you are manually assigning ownership in an HA pair that is not being initialized and does not have only DS460C shelves, use option 1.
- If you are initializing an HA pair that has only DS460C shelves, use option 2 to manually assign ownership for the root drives.

Option 1: Most HA pairs

For an HA pair that is not being initialized and does not have only DS460C shelves, use this procedure to manually assigning ownership.

About this task

- The disks you are assigning ownership for must be in a shelf that is physically cabled to the node you are assigning ownership to.
- If you are using disks in a local tier (aggregate):
 - Disks must be owned by a node before they can be used in a local tier (aggregate).
 - You cannot reassign ownership of a disk that is in use in a local tier (aggregate).

Steps

1. Use the CLI to display all unowned disks:

```
storage disk show -container-type unassigned
```

2. Assign each disk:

```
storage disk assign -disk disk_name -owner owner_name
```

You can use the wildcard character to assign more than one disk at once. If you are reassigning a spare disk that is already owned by a different node, you must use the “-force” option.

Option 2: An HA pair with only DS460C shelves

For an HA pair that you are initializing and that only has DS460C shelves, use this procedure to manually assign ownership for the root drives.

About this task

- When you initialize an HA pair that has only DS460C shelves, you must manually assign the root drives to conform to the half-drawer policy.

After HA pair initialization (boot up), automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the root drives) and any drives added in the future, such as replacing failed disks, responding to a “low spares” message, or adding capacity.

Learn about the half-drawer policy in the topic [About automatic assignment of disk ownership](#).

- RAID needs a minimum of 10 drives for each HA pair (5 for each node) for any greater than 8TB NL-SAS drives in a DS460C shelf.

Steps

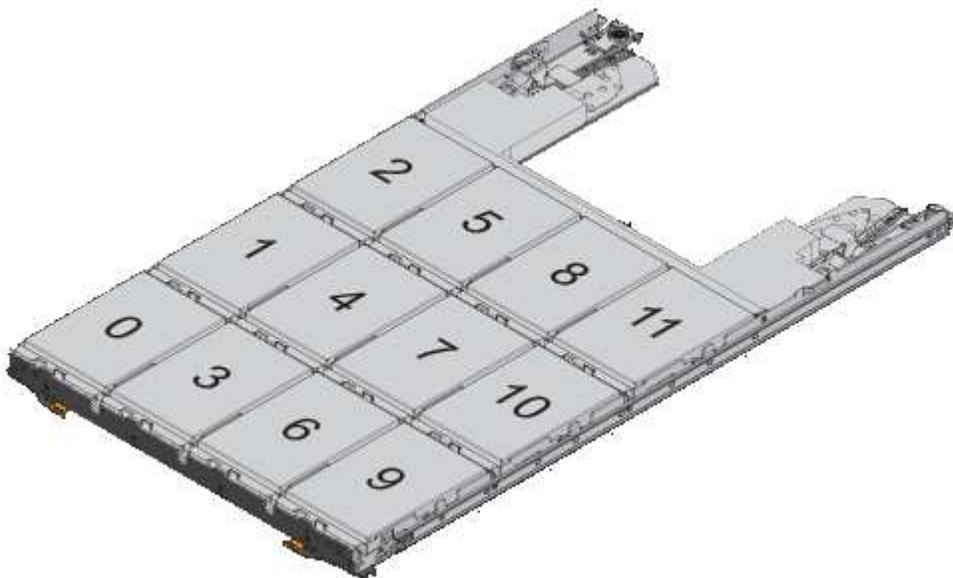
1. If your DS460C shelves are not fully populated, complete the following substeps; otherwise, go to the next step.
 - a. First, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.

- b. For the remaining drives, evenly distribute them across each drawer.

Fill drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

The following illustration shows the drive bay numbering and locations in a DS460C drawer.



2. Log into the clustershell using the node-management LIF or cluster-management LIF.

3. Manually assign the root drives in each drawer to conform to the half-drawer policy using the following substeps:

The half-drawer policy has you assign the left half of a drawer's drives (bays 0 to 5) to node A, and the right half of a drawer's drives (bays 6 to 11) to node B.

- a. Display all unowned disks: `storage disk show -container-type unassigned``
- b. Assign the root disks: `storage disk assign -disk disk_name -owner owner_name`

You can use the wildcard character to assign more than one disk at a time.

Manually assign ownership of partitioned disks

You can manually assign the ownership of the container disk or the partitions on Advanced Drive Partitioning (ADP) systems. If you are initializing an HA pair that only has DS460C shelves, you must manually assign ownership for the container drives that will include root partitions.

About this task

- The type of storage system you have determines which method of ADP is supported, root-data (RD) or root-data-data (RD2).

FAS storage systems use RD and AFF storage systems use RD2.

- If you are manually assigning ownership in an HA pair that is not being initialized and does not have only DS460C shelves, use option 1 to manually assign disks with root-data (RD) partitioning or use option 2 to manually assign disks with root-data-data (RD2) partitioning.
- If you are initializing an HA pair that has only DS460C shelves, use option 3 to manually assign ownership for the container drives that have the root partition.

Option 1: Manually assign disks with root-data (RD) partitioning

For root-data partitioning, there are three owned entities (the container disk and the two partitions) collectively owned by the HA pair.

About this task

- The container disk and the two partitions do not all need to be owned by the same node in the HA pair as long as they are all owned by one of the nodes in the HA pair. However, when you use a partition in a local tier (aggregate), it must be owned by the same node that owns the local tier.
- If a container disk fails in a half-populated shelf and is replaced, you might need to manually assign disk ownership because ONTAP does not always auto-assign ownership in this case.
- After the container disk is assigned, ONTAP's software automatically handles any partitioning and partition assignments that are required.

Steps

1. Use the CLI to display the current ownership for the partitioned disk:

```
storage disk show -disk disk_name -partition-ownership
```

2. Set the CLI privilege level to advanced:

```
set -privilege advanced
```

3. Enter the appropriate command, depending on which ownership entity you want to assign ownership for:

If any of the ownership entities are already owned, then you must include the “-force” option.

If you want to assign ownership for the...	Use this command...
Container disk	<pre>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i></pre>
Data partition	<pre>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data true</pre>
Root partition	<pre>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -root true</pre>

Option 2: Manually assign disks with root-data-data (RD2) partitioning

For root-data-data partitioning, there are four owned entities (the container disk and the three partitions) collectively owned by the HA pair. Root-data-data partitioning creates one small partition as the root partition and two larger, equally sized partitions for data.

About this task

- Parameters must be used with the `disk assign` command to assign the proper partition of a root-data-data partitioned disk. You cannot use these parameters with disks that are part of a storage pool. The default value is “false”.
 - The `-data1 true` parameter assigns the “data1” partition of a root-data1-data2 partitioned disk.
 - The `-data2 true` parameter assigns the “data2” partition of a root-data1-data2 partitioned disk.
- If a container disk fails in a half-populated shelf and is replaced, you might need to manually assign disk ownership because ONTAP does not always auto-assign ownership in this case.
- After the container disk is assigned, ONTAP’s software automatically handles any partitioning and partition assignments that are required.

Steps

1. Use the CLI to display the current ownership for the partitioned disk:

```
storage disk show -disk disk_name -partition-ownership
```

2. Set the CLI privilege level to advanced:

```
set -privilege advanced
```

3. Enter the appropriate command, depending on which ownership entity you want to assign ownership for:

If any of the ownership entities are already owned, then you must include the “-force” option.

If you want to assign ownership for the...	Use this command...
Container disk	<code>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i></code>
Data1 partition	<code>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data1 true</code>
Data2 partition	<code>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data2 true</code>
Root partition	<code>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -root true</code>

Option 3: Manually assign DS460C container drives that have the root partition

If you are initializing an HA pair that has only DS460C shelves, you must manually assign ownership for the container drives that have the root partition by conforming to the half-drawer policy.

About this task

- When you initialize an HA pair that has only DS460C shelves, the ADP boot menu (available with ONTAP 9.2 and later) options 9a and 9b do not support automatic drive ownership assignment. You must manually assign the container drives that have the root partition by conforming to the half-drawer policy.

After HA pair initialization (boot up), automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the container drives that have the root partition) and any drives added in the future, such as replacing failed drives, responding to a “low spares” message, or adding capacity.

- Learn about the half-drawer policy in the topic [About automatic assignment of disk ownership](#).

Steps

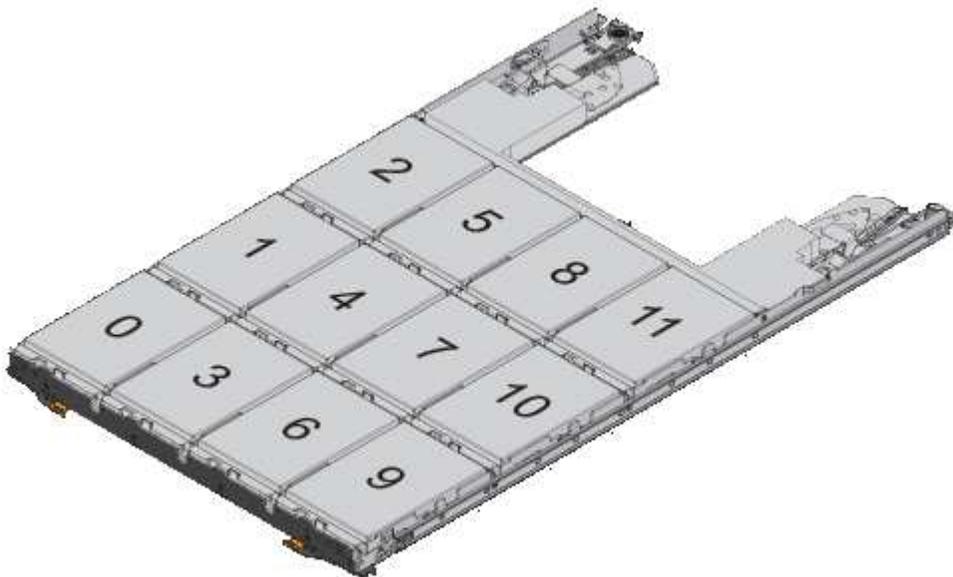
1. If your DS460C shelves are not fully populated, complete the following substeps; otherwise, go to the next step.
 - a. First, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.

- b. For the remaining drives, evenly distribute them across each drawer.

Fill drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

The following illustration shows the drive bay numbering and locations in a DS460C drawer.



2. Log into the clustershell using the node-management LIF or cluster-management LIF.
3. For each drawer, manually assign the container drives that have the root partition by conforming to


```

Root Physical
Disk          Type      RPM Checksum  Usable
Usable      Size
-----
1.0.0        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.1        BSAS      7200 block    753.8GB
73.89GB 828.0GB
1.0.5        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.6        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.10       BSAS      7200 block    753.8GB
0B 828.0GB
1.0.11       BSAS      7200 block    753.8GB
0B 828.0GB

Original Owner: cluster1-02
Pool0
Partitioned Spares

Local
Local
Data
Root Physical
Disk          Type      RPM Checksum  Usable
Usable      Size
-----
1.0.2        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.3        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.4        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.7        BSAS      7200 block    753.8GB
0B 828.0GB
1.0.8        BSAS      7200 block    753.8GB
73.89GB 828.0GB
1.0.9        BSAS      7200 block    753.8GB
0B 828.0GB
12 entries were displayed.

```

2. Enter the advanced privilege level:

```
set advanced
```


Learn about [advanced disk partitioning](#).

Steps

All commands are input at the cluster shell.

1. View the current ownership of the data partitions:

```
storage aggregate show-spare-disks -original-owner passive_node_name -fields local-usable-data1-size, local-usable-data2-size
```

The output shows that half of the data partitions are owned by one node and half are owned by the other node. All of the data partitions should be spare.

2. Enter the advanced privilege level:

```
set advanced
```

3. For each data1 partition owned by the node that will be the passive node, assign it to the active node:

```
storage disk assign -force -data1 -owner active_node_name -disk disk_name
```

You do not need to include the partition as part of the disk name

4. For each data2 partition owned by the node that will be the passive node, assign it to the active node:

```
storage disk assign -force -data2 -owner active_node_name -disk disk_name
```

You do not need to include the partition as part of the disk name

5. Confirm that all of the partitions are assigned to the active node:

```
storage aggregate show-spare-disks
```

```
cluster1::*> storage aggregate show-spare-disks

Original Owner: cluster1-01
Pool0
  Partitioned Spares

Local
Local
Root Physical
Disk          Type      RPM  Checksum  Usable
Usable      Size
-----
1.0.0      BSAS    7200  block    753.8GB
0B 828.0GB
1.0.1      BSAS    7200  block    753.8GB
73.89GB 828.0GB
1.0.2      BSAS    7200  block    753.8GB
```

```

0B 828.0GB
  1.0.3          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.4          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.5          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.6          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.7          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.8          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.9          BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.10         BSAS 7200 block 753.8GB
0B 828.0GB
  1.0.11         BSAS 7200 block 753.8GB
0B 828.0GB

```

Original Owner: cluster1-02

Pool0

Partitioned Spares

```

Local
Local
Root Physical
Disk
Usable Size
-----
1.0.8          BSAS 7200 block 0B
73.89GB 828.0GB
13 entries were displayed.

```

Note that cluster1-02 still owns a spare root partition.

6. Return to administrative privilege:

```
set admin
```

7. Create your data aggregate, leaving at least one data partition as spare:

```
storage aggregate create new_aggr_name -diskcount number_of_partitions -node
active_node_name
```

The data aggregate is created and is owned by the active node.

- Alternatively, you can use ONTAP's recommend aggregate layout which includes best practices for RAID group layout and spare counts:

```
storage aggregate auto-provision
```

Remove ownership from a disk

ONTAP writes disk ownership information to the disk. Before you remove a spare disk or its shelf from a node, you should remove its ownership information so that it can be properly integrated into another node.



If the disk is partitioned for root-data partitioning and you are running ONTAP 9.10.1 or later, contact NetApp Technical Support for assistance in removing ownership. For more information see the [Knowledge Base article: Failed to remove the owner of disk](#).

What you'll need

The disk you want to remove ownership from must meet the following requirements:

- It must be a spare disk.

You cannot remove ownership from a disk that is being used in an local tier (aggregate).

- It cannot be in the maintenance center.
- It cannot be undergoing sanitization.
- It cannot have failed.

It is not necessary to remove ownership from a failed disk.

About this task

If you have automatic disk assignment enabled, ONTAP could automatically reassign ownership before you remove the disk from the node. For this reason, you disable the automatic ownership assignment until the disk is removed, and then you re-enable it.

Steps

- If disk ownership automatic assignment is on, use the CLI to turn it off:

```
storage disk option modify -node node_name -autoassign off
```

- If needed, repeat the previous step for the node's HA partner.
- Remove the software ownership information from the disk:

```
storage disk removeowner disk_name
```

To remove ownership information from multiple disks, use a comma-separated list.

Example:

```
storage disk removeowner sys1:0a.23,sys1:0a.24,sys1:0a.25
```

4. If the disk is partitioned for root-data partitioning and you are running ONTAP 9.9.1 or earlier, remove ownership from the partitions:

```
storage disk removeowner -disk disk_name -root true
```

```
storage disk removeowner -disk disk_name -data true
```

Both partitions are no longer owned by any node.

5. If you previously turned off automatic assignment of disk ownership, turn it on after the disk has been removed or reassigned:

```
storage disk option modify -node node_name -autoassign on
```

6. If needed, repeat the previous step for the node's HA partner.

Remove a failed disk

A disk that has completely failed is no longer counted by ONTAP as a usable disk, and you can immediately disconnect the disk from the disk shelf. However, you should leave a partially failed disk connected long enough for the Rapid RAID Recovery process to complete.

About this task

If you are removing a disk because it has failed or because it is producing excessive error messages, you should not use the disk again in this or any other storage system.

Steps

1. Use the CLI to find the disk ID of the failed disk:

```
storage disk show -broken
```

If the disk does not appear in the list of failed disks, it might have partially failed, with a Rapid RAID Recovery in process. In this case, you should wait until the disk is present in the list of failed disks (which means that the Rapid RAID Recovery process is complete) before removing the disk.

2. Determine the physical location of the disk you want to remove:

```
storage disk set-led -action on -disk disk_name 2
```

The fault LED on the face of the disk is lit.

3. Remove the disk from the disk shelf, following the instructions in the hardware guide for your disk shelf model.

Disk sanitization

Disk sanitization overview

Disk sanitization is the process of physically obliterating data by overwriting disks or SSDs with specified byte patterns or random data so that recovery of the original data

becomes impossible. Using the sanitization process ensures that no one can recover the data on the disks.

This functionality is available through the nodeshell in all ONTAP 9 releases, and starting with ONTAP 9.6 in maintenance mode.

The disk sanitization process uses three successive default or user-specified byte overwrite patterns for up to seven cycles per operation. The random overwrite pattern is repeated for each cycle.

Depending on the disk capacity, the patterns, and the number of cycles, the process can take several hours. Sanitization runs in the background. You can start, stop, and display the status of the sanitization process. The sanitization process contains two phases: the "Formatting phase" and the "Pattern overwrite phase".

Formatting phase

The operation performed for the formatting phase depends on the class of disk being sanitized, as shown in the following table:

Disk class	Formatting phase operation
Capacity HDDs	Skipped
Performance HDDs	SCSI format operation
SSDs	SCSI sanitize operation

Pattern overwrite phase

The specified overwrite patterns are repeated for the specified number of cycles.

When the sanitization process is complete, the specified disks are in a sanitized state. They are not returned to spare status automatically. You must return the sanitized disks to the spare pool before the newly sanitized disks are available to be added to another aggregate.

When disk sanitization cannot be performed

Disk sanitization is not supported for all disk types. In addition, there are circumstances in which disk sanitization cannot be performed.

- It is not supported on all SSD part numbers.

For information about which SSD part numbers support disk sanitization, see the [Hardware Universe](#).

- It is not supported in takeover mode for systems in an HA pair.
- It cannot be performed on disks that were failed due to readability or writability problems.
- It does not perform its formatting phase on ATA drives.
- If you are using the random pattern, it cannot be performed on more than 100 disks at one time.
- It is not supported on array LUNs.
- If you sanitize both SES disks in the same ESH shelf at the same time, you see errors on the console about access to that shelf, and shelf warnings are not reported for the duration of the sanitization.

However, data access to that shelf is not interrupted.

What happens if disk sanitization is interrupted

If disk sanitization is interrupted by user intervention or an unexpected event such as a power outage, ONTAP takes action to return the disks that were being sanitized to a known state, but you must also take action before the sanitization process can finish.

Disk sanitization is a long-running operation. If the sanitization process is interrupted by power failure, system panic, or manual intervention, the sanitization process must be repeated from the beginning. The disk is not designated as sanitized.

If the formatting phase of disk sanitization is interrupted, ONTAP must recover any disks that were corrupted by the interruption. After a system reboot and once every hour, ONTAP checks for any sanitization target disk that did not complete the formatting phase of its sanitization. If any such disks are found, ONTAP recovers them. The recovery method depends on the type of the disk. After a disk is recovered, you can rerun the sanitization process on that disk; for HDDs, you can use the `-s` option to specify that the formatting phase is not repeated again.

Tips for creating and backing up local tiers (aggregates) containing data to be sanitized

If you are creating or backing up local tiers (aggregates) to contain data that might need to be sanitized, following some simple guidelines will reduce the time it takes to sanitize your data.

- Make sure your local tiers containing sensitive data are not larger than they need to be.

If they are larger than needed, sanitization requires more time, disk space, and bandwidth.

- When you back up local tiers containing sensitive data, avoid backing them up to local tier that also contain large amounts of nonsensitive data.

This reduces the resources required to move nonsensitive data before sanitizing sensitive data.

Sanitize a disk

Sanitizing a disk allows you to remove data from a disk or a set of disks on decommissioned or inoperable systems so that the data can never be recovered.

Two methods are available to sanitize disks using the CLI:

Sanitize a disk with “maintenance mode” commands (ONTAP 9.6 and later releases)

Beginning with ONTAP 9.6, you can perform disk sanitization in maintenance mode.

Before you begin

- The disks cannot be self-encrypting disks (SED).

You must use the `storage encryption disk sanitize` command to sanitize an SED.

[Encryption of data at rest](#)

Steps

1. Boot into maintenance mode.

- a. Exit the current shell by entering `halt`.

The LOADER prompt is displayed.

- b. Enter maintenance mode by entering `boot_ontap maint`.

After some information is displayed, the maintenance mode prompt is displayed.

2. If the disks you want to sanitize are partitioned, unpartition each disk:



The command to unpartition a disk is only available at the diag level and should be performed only under NetApp Support supervision. It is highly recommended that you contact NetApp Support before you proceed. You can also refer to the Knowledge Base article [How to unpartition a spare drive in ONTAP](#)

```
disk unpartition disk_name
```

3. Sanitize the specified disks:

```
disk sanitize start [-p pattern1|-r [-p pattern2|-r [-p pattern3|-r]]] [-c cycle_count] disk_list
```



Do not turn off power to the node, disrupt the storage connectivity, or remove target disks while sanitizing. If sanitizing is interrupted during the formatting phase, the formatting phase must be restarted and allowed to finish before the disks are sanitized and ready to be returned to the spare pool. If you need to abort the sanitization process, you can do so by using the `disk sanitize abort` command. If the specified disks are undergoing the formatting phase of sanitization, the abort does not occur until the phase is complete.

`-p pattern1 -p pattern2 -p pattern3` specifies a cycle of one to three user-defined hex byte overwrite patterns that can be applied in succession to the disks being sanitized. The default pattern is three passes, using 0x55 for the first pass, 0xaa for the second pass, and 0x3c for the third pass.

`-r` replaces a patterned overwrite with a random overwrite for any or all of the passes.

`-c cycle_count` specifies the number of times that the specified overwrite patterns are applied. The default value is one cycle. The maximum value is seven cycles.

disk_list specifies a space-separated list of the IDs of the spare disks to be sanitized.

4. If desired, check the status of the disk sanitization process:

```
disk sanitize status [disk_list]
```

5. After the sanitization process is complete, return the disks to spare status for each disk:

```
disk sanitize release disk_name
```

6. Exit maintenance mode.

Sanitize a disk with “nodeshell” commands (all ONTAP 9 releases)

For all versions of ONTAP 9, when disk sanitization is enabled using nodeshell commands, some low-level ONTAP commands are disabled. After disk sanitization is enabled on a node, it cannot be disabled.

Before you begin

- The disks must be spare disks; they must be owned by a node, but not used in a local tier (aggregate).

If the disks are partitioned, neither partition can be in use in a local tier (aggregate).

- The disks cannot be self-encrypting disks (SED).

You must use the `storage encryption disk sanitize` command to sanitize an SED.

Encryption of data at rest

- The disks cannot be part of a storage pool.

Steps

1. If the disks you want to sanitize are partitioned, unpartition each disk:



The command to unpartition a disk is only available at the diag level and should be performed only under NetApp Support supervision. **It is highly recommended that you contact NetApp Support before you proceed.** You can also refer to the Knowledge Base article [How to unpartition a spare drive in ONTAP](#).

```
disk unpartition disk_name
```

2. Enter the nodeshell for the node that owns the disks you want to sanitize:

```
system node run -node node_name
```

3. Enable disk sanitization:

```
options licensed_feature.disk_sanitization.enable on
```

You are asked to confirm the command because it is irreversible.

4. Switch to the nodeshell advanced privilege level:

```
priv set advanced
```

5. Sanitize the specified disks:

```
disk sanitize start [-p pattern1|-r [-p pattern2|-r [-p pattern3|-r]]] [-c  
cycle_count] disk_list
```



Do not turn off power to the node, disrupt the storage connectivity, or remove target disks while sanitizing. If sanitizing is interrupted during the formatting phase, the formatting phase must be restarted and allowed to finish before the disks are sanitized and ready to be returned to the spare pool. If you need to abort the sanitization process, you can do so by using the `disk sanitize abort` command. If the specified disks are undergoing the formatting phase of sanitization, the abort does not occur until the phase is complete.

`-p pattern1 -p pattern2 -p pattern3` specifies a cycle of one to three user-defined hex byte overwrite patterns that can be applied in succession to the disks being sanitized. The default pattern is three passes, using 0x55 for the first pass, 0xaa for the second pass, and 0x3c for the third pass.

`-r` replaces a patterned overwrite with a random overwrite for any or all of the passes.

`-c cycle_count` specifies the number of times that the specified overwrite patterns are applied.

The default value is one cycle. The maximum value is seven cycles.

`disk_list` specifies a space-separated list of the IDs of the spare disks to be sanitized.

6. If you want to check the status of the disk sanitization process:

```
disk sanitize status [disk_list]
```

7. After the sanitization process is complete, return the disks to spare status:

```
disk sanitize release disk_name
```

8. Return to the nodeshell admin privilege level:

```
priv set admin
```

9. Return to the ONTAP CLI:

```
exit
```

10. Determine whether all of the disks were returned to spare status:

```
storage aggregate show-spare-disks
```

If...	Then...
All of the sanitized disks are listed as spares	You are done. The disks are sanitized and in spare status.

Some of the sanitized disks are not listed as spares

Complete the following steps:

a. Enter advanced privilege mode:

```
set -privilege advanced
```

b. Assign the unassigned sanitized disks to the appropriate node for each disk:

```
storage disk assign -disk disk_name -owner  
node_name
```

c. Return the disks to spare status for each disk:

```
storage disk unfail -disk disk_name -s -q
```

d. Return to administrative mode:

```
set -privilege admin
```

Result

The specified disks are sanitized and designated as hot spares. The serial numbers of the sanitized disks are written to `/etc/log/sanitized_disks`.

The specified disks' sanitization logs, which show what was completed on each disk, are written to `/mroot/etc/log/sanitization.log`.

Commands for managing disks

You can use the `storage disk` and `storage aggregate` commands to manage your disks.

If you want to...	Use this command...
Display a list of spare disks, including partitioned disks, by owner	<code>storage aggregate show-spare-disks</code>
Display the disk RAID type, current usage, and RAID group by aggregate	<code>storage aggregate show-status</code>
Display the RAID type, current usage, aggregate, and RAID group, including spares, for physical disks	<code>storage disk show -raid</code>
Display a list of failed disks	<code>storage disk show -broken</code>
Display the pre-cluster (nodescope) drive name for a disk	<code>storage disk show -primary-paths (advanced)</code>

Illuminate the LED for a particular disk or shelf	<code>storage disk set-led</code>
Display the checksum type for a specific disk	<code>storage disk show -fields checksum-compatibility</code>
Display the checksum type for all spare disks	<code>storage disk show -fields checksum-compatibility -container-type spare</code>
Display disk connectivity and placement information	<code>storage disk show -fields disk,primary-port,secondary-name,secondary-port,shelf,bay</code>
Display the pre-cluster disk names for specific disks	<code>storage disk show -disk diskname -fields diskpathnames</code>
Display the list of disks in the maintenance center	<code>storage disk show -maintenance</code>
Display SSD wear life	<code>storage disk show -ssd-wear</code>
Unpartition a shared disk	<code>storage disk unpartition (available at diagnostic level)</code>
Zero all non-zeroed disks	<code>storage disk zerospares</code>
Stop an ongoing sanitization process on one or more specified disks	<code>system node run -node nodename -command disk sanitize</code>
Display storage encryption disk information	<code>storage encryption disk show</code>
Retrieve authentication keys from all linked key management servers	<code>security key-manager restore</code>

Related information

[ONTAP 9 Commands](#)

Commands for displaying space usage information

You use the `storage aggregate` and `volume` commands to see how space is being used in your aggregates and volumes and their Snapshot copies.

To display information about...	Use this command...
---------------------------------	---------------------

Aggregates, including details about used and available space percentages, Snapshot reserve size, and other space usage information	<code>storage aggregate show</code> <code>storage aggregate show-space -fields snap-size-total,used-including-snapshot-reserve</code>
How disks and RAID groups are used in an aggregate, and RAID status	<code>storage aggregate show-status</code>
The amount of disk space that would be reclaimed if you deleted a specific Snapshot copy	<code>volume snapshot compute-reclaimable</code>
The amount of space used by a volume	<code>volume show -fields size,used,available,percent-used</code> <code>volume show-space</code>
The amount of space used by a volume in the containing aggregate	<code>volume show-footprint</code>

Related information

[ONTAP 9 Commands](#)

Commands for displaying information about storage shelves

You use the `storage shelf show` command to display configuration and error information for your disk shelves.

If you want to display...	Use this command...
General information about shelf configuration and hardware status	<code>storage shelf show</code>
Detailed information for a specific shelf, including stack ID	<code>storage shelf show -shelf</code>
Unresolved, customer actionable, errors by shelf	<code>storage shelf show -errors</code>
Bay information	<code>storage shelf show -bay</code>
Connectivity information	<code>storage shelf show -connectivity</code>
Cooling information, including temperature sensors and cooling fans	<code>storage shelf show -cooling</code>
Information about I/O modules	<code>storage shelf show -module</code>

If you want to display...	Use this command...
Port information	<code>storage shelf show -port</code>
Power information, including PSUs (power supply units), current sensors, and voltage sensors	<code>storage shelf show -power</code>

Related information

[ONTAP 9 Commands](#)

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.