



ONTAP® 9

Disks and Aggregates Power Guide

February 2019 | 215-11204_LO
doccomments@netapp.com

Updated for ONTAP 9.5

 **NetApp**®

Contents

Deciding whether to use this guide	6
Aggregate creation workflow	7
Deciding which aggregate creation method to use	7
Creating aggregates with auto-provision	8
Default RAID policies for aggregates	9
How to determine the number of disks or disk partitions required for an aggregate	9
Creating aggregates manually	10
Aggregate expansion workflow	12
Adding disks to a node	13
Manually assigning disk ownership	13
Expanding aggregates	14
Managing aggregates	17
RAID protection levels for disks	17
How to determine the number of disks or disk partitions required for an aggregate	17
Correcting misaligned spare partitions	18
Determining drive and RAID group information for an aggregate	19
Relocating aggregate ownership within an HA pair	20
Relocating aggregate ownership	20
Commands for aggregate relocation	22
Assigning aggregates to SVMs	22
How to determine space usage in an aggregate	23
Determining which volumes reside on an aggregate	24
How you can determine and control a volume's space usage in the aggregate	24
Methods to create space in an aggregate	25
Commands for managing aggregates	26
Working with Flash Pool aggregates	27
Flash Pool caching policies and SSD partitioning	27
How Flash Pool aggregate caching policies work	27
How Flash Pool SSD partitioning works for Flash Pool aggregates using storage pools	27
Determining Flash Pool candidacy and optimal cache size	29
Creating a Flash Pool aggregate using physical SSDs	31
Creating a Flash Pool aggregate using SSD storage pools	32
Determining whether a Flash Pool aggregate is using an SSD storage pool	32
Creating an SSD storage pool	32
Creating a Flash Pool aggregate using SSD storage pool allocation units	33
Determining the impact to cache size of adding SSDs to an SSD storage pool	34

Adding SSDs to an SSD storage pool	35
Commands for managing SSD storage pools	35
Determining whether to modify the caching policy of Flash Pool aggregates	36
Modifying caching policies of Flash Pool aggregates	36
Setting the cache-retention policy for Flash Pool aggregates	37
Managing storage tiers by using FabricPool	39
Benefits of storage tiers by using FabricPool	39
Considerations and requirements for using FabricPool	39
Understanding FabricPool tiering policies	42
FabricPool management workflow	44
Configuring FabricPool	44
Preparing for FabricPool configuration	45
Setting up an object store as the cloud tier for FabricPool	47
Attaching the cloud tier to an aggregate	50
Managing FabricPool	51
Determining how much data in a volume is inactive by using inactive data reporting	51
Adding or moving volumes to FabricPool as needed	53
Monitoring the space utilization for FabricPool	55
Managing storage tiering by modifying a volume's tiering policy or tiering minimum cooling period	56
Commands for managing aggregates with FabricPool	58
Managing disks	60
When you need to update the Disk Qualification Package	60
How hot spare disks work	60
How low spare warnings can help you manage your spare disks	60
Displaying disk and partition ownership	61
Manually assigning ownership of partitioned disks	61
Root-data partitioning	62
Root-data-data partitioning	62
Additional root-data partitioning management options	63
Configuring autoassignment of disk ownership	64
Which disk autoassignment policy to use	64
Removing a failed disk	65
Removing ownership from a disk	65
Sanitizing a disk	66
Setting up an active-passive configuration on nodes using root-data partitioning	68
Commands for managing disks	70
Commands for displaying space usage information	71
Commands for displaying information about storage shelves	71
Managing RAID groups	73
Converting from RAID-DP to RAID-TEC	73
Converting RAID-TEC to RAID-DP	73
Considerations for sizing RAID groups	74
Customizing the size of your RAID groups	74

Appendix: Mirrored and unmirrored aggregates	76
How unmirrored aggregates work	76
How mirrored aggregates work	77
Where to find additional information	79
Copyright	80
Trademark	81
How to send comments about documentation and receive update notifications	82
Index	83

Deciding whether to use the Disks and Aggregates Power Guide

This guide describes how to manage ONTAP physical storage. It shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies.

You should use this guide under the following circumstances:

- You want to use the command-line interface (CLI), not OnCommand System Manager or an automated scripting tool.
- You want to use best practices, not explore every available option.
- You have a MetroCluster configuration and you are following the procedures in the MetroCluster documentation for initial configuration and guidelines for aggregates and disk management.

If you want to use OnCommand System Manager to manage disks and aggregates, you should choose the following documentation:

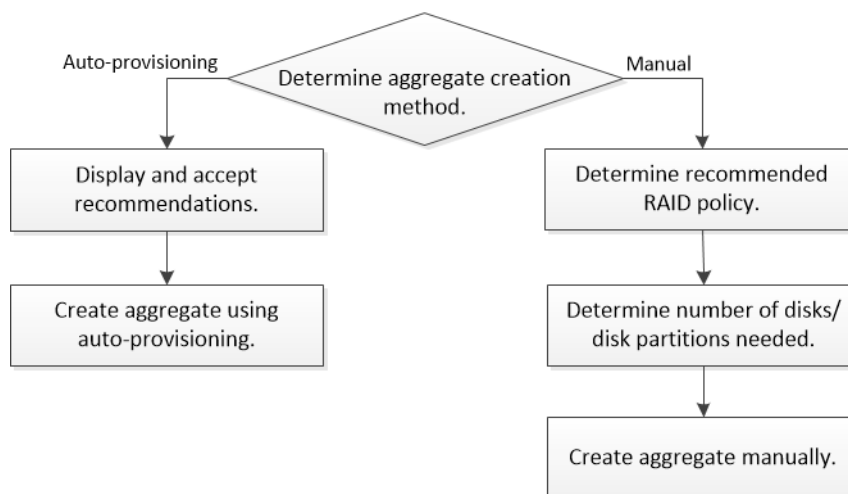
- *Cluster management using System Manager*

If you require additional configuration or conceptual information, you should choose among the following documentation:

- Conceptual background for disk and aggregate management
ONTAP concepts
- FlexVol volumes, FlexClone technology, and storage efficiency features
Logical storage management
- NAS file access
 - *NFS management*
 - *SMB/CIFS management*
- SAN host provisioning
 - *SAN administration*
- MetroCluster documentation
 - *Fabric-attached MetroCluster installation and configuration*
 - *Stretch MetroCluster installation and configuration*
 - *MetroCluster IP installation and configuration*
- Command reference
ONTAP 9 commands
- Automation of management tasks
NetApp Documentation: OnCommand Workflow Automation (current releases)

Aggregate creation workflow

Creating aggregates provides storage to volumes on your system. Beginning in ONTAP 9.2, you can let ONTAP recommend aggregate configurations for your system (auto-provision). If the auto-provision method is not available or appropriate in your environment, you can configure aggregates manually.



Deciding which aggregate creation method to use

Although aggregate creation with auto-provision is a best practice in ONTAP 9.2 and later, you must determine whether it is supported in your environment. If it is not, you must make decisions about RAID policy and disk configuration, and then create the aggregates manually.

When you create an aggregate using the `storage aggregate auto-provision` command, ONTAP analyzes available spare disks in the cluster and generates a recommendation about how spare disks should be used to create aggregates according to best practices. ONTAP displays the summary of recommended aggregates including their names and usable size, and then prompts you to decide whether the aggregates should be created as recommended.

In many cases, the recommended aggregate layout in the auto-provision display will be optimal for your environment. However, if your cluster is running ONTAP 9.1 or earlier, or your environment includes the following configurations, you must use the manual aggregate configuration method.

- Aggregates using third-party array LUNs
- Virtual disks with Cloud Volumes ONTAP or ONTAP Select
- MetroCluster
- SyncMirror
- MSATA disks
- FlashPool aggregates
- Multiple disk types or sizes are connected to the node

In addition, if any of the following disk conditions are present, they must be addressed before using the auto-provision method:

- Missing disks
- Fluctuation in spare disk numbers
- Unassigned disks
- Non-zeroed spares
- Disks undergoing maintenance testing

The `storage aggregate auto-provision man` page contains more information about these requirements.

Related information

[ONTAP 9 commands](#)

Creating aggregates with auto-provision

If the auto-provision method is appropriate in your environment, you run the `storage aggregate auto-provision` to generate aggregate layout recommendations. You can then create aggregates after reviewing and approving ONTAP recommendations.

Before you begin

ONTAP 9.2 or later must be running on your cluster.

About this task

The default summary generated with the `storage aggregate auto-provision` command lists the recommended aggregates to be created, including names and usable size. You can view the list and determine whether you want to create the recommended aggregates when prompted.

You can also display a detailed summary by using the `-verbose` option, which displays the following reports:

- Per node summary of new aggregates to create, discovered spares, and remaining spare disks and partitions after aggregate creation
- New data aggregates to create with counts of disks and partitions to be used
- RAID group layout showing how spare disks and partitions will be used in new data aggregates to be created
- Details about spare disks and partitions remaining after aggregate creation

If you are familiar with the auto-provision method and your environment is correctly prepared, you can use the `-skip-confirmation` option to create the recommended aggregate without display and confirmation. The `storage aggregate auto-provision` command is not affected by the CLI session `-confirmations` setting.

The `storage aggregate auto-provision man` page contains more information about the aggregate layout recommendations.

Steps

1. Run the `storage aggregate auto-provision` command with the desired display options.
 - no options: Display standard summary
 - `-verbose` option: Display detailed summary

- `-skip-confirmation` option: Create recommended aggregates without display or confirmation
2. After reviewing the display of recommended aggregates, respond to the prompt to create the recommended aggregates.

Example

```
Do you want to create recommended aggregates? {y|n}:y
Info: Creating node1_SSD_1 ...
      Creating node2_SSD_1 ...
```

Related information

[ONTAP 9 commands](#)

Default RAID policies for aggregates

Either RAID-DP or RAID-TEC is the default RAID policy for all new aggregates. The RAID policy determines the parity protection you have in the event of a disk failure.

RAID-DP provides double-parity protection in the event of a single or double disk failure. RAID-DP is the default RAID policy for the following aggregate types:

- All flash aggregates
- Flash Pool aggregates
- Performance hard disk drive (HDD) aggregates

Beginning with ONTAP 9.0, a new RAID policy called RAID-TEC is available. RAID-TEC is supported on all disk types and all platforms, including AFF. Aggregates that contain larger disks have a higher possibility of concurrent disk failures. RAID-TEC helps to mitigate this risk by providing triple-parity protection so that your data can survive up to three simultaneous disk failures. RAID-TEC is the default RAID policy for capacity HDD aggregates with disks that are 6 TB or larger.

How to determine the number of disks or disk partitions required for an aggregate

You must have enough disks or disk partitions in your aggregate to meet system and business requirements. You should also have the recommended number of hot spare disks or hot spare disk partitions to minimize the potential of data loss.

Root-data partitioning is enabled by default on certain configurations. Systems with root-data partitioning enabled use disk partitions to create aggregates. Systems that do not have root-data partitioning enabled use unpartitioned disks.

You must have enough disks or disk partitions to meet the minimum number required for your RAID policy and enough to meet your minimum capacity requirements.

Note: In ONTAP, the usable space of the drive is less than the physical capacity of the drive. You can find the usable space of a specific drive and the minimum number of disks or disk partitions required for each RAID policy in *Hardware Universe*. You can also use the `storage aggregate show-spare-disks` command to find the usable space of a specific disk.

In addition to the number of disks or disk partitions necessary to create your RAID group and meet your capacity requirements, you should also have the minimum number of hot spare disks or hot spare disk partitions recommended for your aggregate:

- For all flash aggregates, you should have a minimum of one hot spare disk or disk partition.
- For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions.
- For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.
- To support the use of the Maintenance Center and to avoid issues caused by multiple concurrent disk failures, you should have a minimum of four hot spares in multi-disk carriers.

Related information

[NetApp Hardware Universe](#)

[NetApp Technical Report 3838: Storage Subsystem Configuration Guide](#)

Creating aggregates manually

Before you create aggregates manually, you should review disk configuration options and simulate creation. Then you can issue the `storage aggregate create` and verify the results.

Before you begin

You must have determined the number of disks and the number of hot spare disks you need in the aggregate.

About this task

If root-data-data partitioning is enabled and you have 24 solid state drives (SSDs) or fewer in your configuration, it is recommended that your data partitions be assigned to different nodes.

The procedure for creating aggregates on systems with root-data partitioning and root-data-data partitioning enabled is the same as the procedure for creating aggregates on systems using unpartitioned disks. If root-data partitioning is enabled on your system, you should use the number of disk partitions for the `-diskcount` option. For root-data-data partitioning, the `-diskcount` option specifies the count of disks to use.

Note: When creating multiple aggregates for use with FlexGroups, aggregates should be as close in size as possible.

The `storage aggregate create man` page contains more information about aggregate creation options and requirements.

Steps

1. View the list of spare disk partitions to verify that you have enough to create your aggregate:

```
storage aggregate show-spare-disks -original-owner node_name
```

Data partitions are displayed under Local Data Usable. A root partition cannot be used as a spare.

2. Simulate the creation of the aggregate:

```
storage aggregate create -aggregate aggregate_name -node node_name -
raidtype raid_dp -diskcount number_of_disks_or_partitions -simulate true
```

3. If any warnings are displayed from the simulated command, adjust the command and repeat the simulation.

4. Create the aggregate:

```
storage aggregate create -aggregate aggr_name -node node_name -raidtype  
raid_dp -diskcount number_of_disks_or_partitions
```

5. Display the aggregate to verify that it was created:

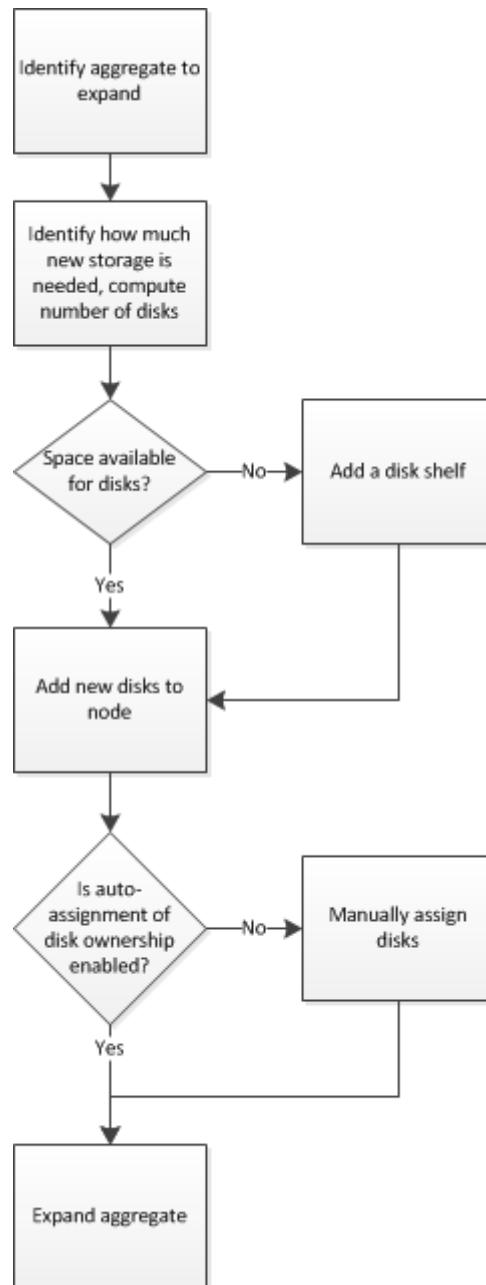
```
storage aggregate show-status aggregate_name
```

Related information

[ONTAP 9 commands](#)

Aggregate expansion workflow

Expanding an aggregate involves identifying the aggregate to expand, determining how much new storage is needed, installing new disks, assigning disk ownership, and creating new a RAID group if needed.



Adding disks to a node

You add disks to a node to increase the number of hot spares, to add space to an aggregate, or to replace disks.

Before you begin

You must have confirmed that your platform model supports the type of disk you want to add.

Steps

1. Check the NetApp Support Site for newer disk and shelf firmware and Disk Qualification Package files.
If your node does not have the latest versions, you must update them before installing the new disk.
2. Install the disks according to the hardware guide for your disk shelf or the hardware and service guide for your platform.

The new disks are not recognized until they are assigned to a node and pool. You can assign the new disks manually, or you can wait for ONTAP to automatically assign the new disks if your node follows the rules for disk autoassignment.

3. After the new disks have all been recognized, verify their addition and their ownership information:

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

You should see the new disks, owned by the correct node and in the correct pool.

4. Optional: Zero the newly added disks:

```
storage disk zerospares
```

Disks that have been used previously in an ONTAP aggregate must be zeroed before they can be added to another aggregate. Zeroing the disks now can prevent delays in case you need to quickly increase the size of an aggregate. The disk zeroing command runs in the background and can take hours to complete, depending on the size of the non-zeroed disks in the node.

Result

The new disks are ready to be added to an aggregate, used to replace an existing disk, or placed onto the list of hot spares.

Related concepts

[When you need to update the Disk Qualification Package](#) on page 60

Manually assigning disk ownership

Disks must be owned by a node before they can be used in an aggregate. If your cluster is not configured to use automatic disk ownership assignment, you must assign ownership manually. You cannot reassign ownership of a disk that is in use in an aggregate.

Steps

1. 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

Expanding aggregates

You can add disks to an aggregate so that it can provide more storage to its associated volumes. The procedure for adding partitioned disks to an aggregate is similar to the procedure for adding unpartitioned disks.

Before you begin

You must know what the RAID group size is for the aggregate you are adding the storage to.

About this task

When you expand an aggregate, you should be aware of whether you are adding partition or unpartitioned disks to the aggregate. When you add unpartitioned drives to an existing aggregate, the size of the existing RAID groups is inherited by the new RAID group, which can affect the number of parity disks required. If an unpartitioned disk is added to a RAID group composed of partitioned disks, the new disk is partitioned, leaving an unused spare partition.

When you provision partitions, you must ensure that you do not leave the node without a drive with both partitions as spare. If you do, and the node experiences a controller disruption, valuable information about the problem (the core file) might not be available to provide to the technical support.

Note: Do not use the `disklist` command to expand your aggregates. This could cause partition misalignment.

Steps

1. Show the available spare storage on the system that owns the aggregate:

```
storage aggregate show-spare-disks -original-owner node_name
```

You can use the `-is-disk-shared` parameter to show only partitioned drives or only unpartitioned drives.

Example

```
c11-s2:> storage aggregate show-spare-disks -original-owner c11-s2 -is-disk-shared true

Original Owner: c11-s2
Pool0
  Shared HDD Spares

Disk          Type      RPM  Checksum      Local Data Usable  Local Root Usable  Physical Size  Status
-----
1.0.1         BSAS     7200 block        753.8GB  73.89GB  828.0GB zeroed
1.0.2         BSAS     7200 block        753.8GB   0B      828.0GB zeroed
1.0.3         BSAS     7200 block        753.8GB   0B      828.0GB zeroed
1.0.4         BSAS     7200 block        753.8GB   0B      828.0GB zeroed
1.0.8         BSAS     7200 block        753.8GB   0B      828.0GB zeroed
1.0.9         BSAS     7200 block        753.8GB   0B      828.0GB zeroed
1.0.10        BSAS     7200 block         0B      73.89GB  828.0GB zeroed
2 entries were displayed.
```

2. Show the current RAID groups for the aggregate:

```
storage aggregate show-status aggr_name
```

Example

```
c11-s2:> storage aggregate show-status -aggregate data_1

Owner Node: c11-s2
Aggregate: data_1 (online, raid_dp) (block checksums)
Plex: /data_1/plex0 (online, normal, active, pool0)
RAID Group /data_1/plex0/rg0 (normal, block checksums)

-----
Position Disk                Pool Type    RPM    Usable Physical
                Size      Size Status
-----
shared  1.0.10                0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.5                  0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.6                  0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.11                 0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.0                  0  BSAS   7200  753.8GB  828.0GB (normal)
5 entries were displayed.
```

3. Simulate adding the storage to the aggregate:

```
storage aggregate add-disks -aggregate aggr_name -diskcount
number_of_disks_or_partitions -simulate true
```

You can see the result of the storage addition without actually provisioning any storage. If any warnings are displayed from the simulated command, you can adjust the command and repeat the simulation.

Example

```
c11-s2:> storage aggregate add-disks data_1 -diskcount 5 -simulate true

Addition of disks would succeed for aggregate "data_1" on node "c11-s2". The
following disks would be used to add to the aggregate: 1.0.2, 1.0.3, 1.0.4, 1.0.8, 1.0.9.
```

4. Add the storage to the aggregate:

```
storage aggregate add-disks -aggregate aggr_name -raidgroup new -
diskcount number_of_disks_or_partitions
```

When creating a Flash Pool aggregate, if you are adding disks with a different checksum than the aggregate, or if you are adding disks to a mixed checksum aggregate, you must use the `-checksumstyle` parameter.

If you are adding disks to a Flash Pool aggregate, you must use the `-disktype` parameter to specify the disk type.

You can use the `-disksize` parameter to specify a size of the disks to add. Only disks with approximately the specified size are selected for addition to the aggregate.

Example

```
c11-s2:> storage aggregate add-disks -aggregate data_1 -raidgroup new -diskcount 5
```

5. Verify that the storage was added successfully:

```
storage aggregate show-status -aggregate aggr_name
```

Example

```
c11-s2:> storage aggregate show-status -aggregate data_1

Owner Node: c11-s2
Aggregate: data_1 (online, raid_dp) (block checksums)
Plex: /data_1/plex0 (online, normal, active, pool0)
RAID Group /data_1/plex0/rg0 (normal, block checksums)

-----
Position Disk                Pool Type    RPM    Usable Physical
                Size      Size Status
-----
shared  1.0.10                0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.5                  0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.6                  0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.11                 0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.0                  0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.2                  0  BSAS   7200  753.8GB  828.0GB (normal)
shared  1.0.3                  0  BSAS   7200  753.8GB  828.0GB (normal)
```

16 | Disks and Aggregates Power Guide

```
shared 1.0.4          0 BSAS 7200 753.8GB 828.0GB (normal)
shared 1.0.8          0 BSAS 7200 753.8GB 828.0GB (normal)
shared 1.0.9          0 BSAS 7200 753.8GB 828.0GB (normal)
10 entries were displayed.
```

6. Verify that the node still has at least one drive with both the root partition and the data partition as spare:

```
storage aggregate show-spare-disks -original-owner node_name
```

Example

```
c11-s2::> storage aggregate show-spare-disks -original-owner c11-s2 -is-disk-shared true
Original Owner: c11-s2
Pool0
  Shared HDD Spares

```

Disk	Type	RPM	Checksum	Local Data Usable	Local Root Usable	Physical Size	Status
1.0.1	BSAS	7200	block	753.8GB	73.89GB	828.0GB	zeroed
1.0.10	BSAS	7200	block	0B	73.89GB	828.0GB	zeroed

```
2 entries were displayed.
```


Managing aggregates

You create and manage your aggregates so that they can provide storage to their associated volumes.

RAID protection levels for disks

ONTAP supports three levels of RAID protection for aggregates. Your level of RAID protection determines the number of parity disks available for data recovery in the event of disk failures.

With RAID protection, if there is a data disk failure in a RAID group, ONTAP can replace the failed disk with a spare disk and use parity data to reconstruct the data of the failed disk.

- **RAID4**
With RAID4 protection, ONTAP can use one spare disk to replace and reconstruct the data from one failed disk within the RAID group.
- **RAID-DP**
With RAID-DP protection, ONTAP can use up to two spare disks to replace and reconstruct the data from up to two simultaneously failed disks within the RAID group.
- **RAID-TEC**
With RAID-TEC protection, ONTAP can use up to three spare disks to replace and reconstruct the data from up to three simultaneously failed disks within the RAID group.

Related information

[NetApp Technical Report 3437: Storage Subsystem Resiliency Guide](#)

How to determine the number of disks or disk partitions required for an aggregate

You must have enough disks or disk partitions in your aggregate to meet system and business requirements. You should also have the recommended number of hot spare disks or hot spare disk partitions to minimize the potential of data loss.

Root-data partitioning is enabled by default on certain configurations. Systems with root-data partitioning enabled use disk partitions to create aggregates. Systems that do not have root-data partitioning enabled use unpartitioned disks.

You must have enough disks or disk partitions to meet the minimum number required for your RAID policy and enough to meet your minimum capacity requirements.

Note: In ONTAP, the usable space of the drive is less than the physical capacity of the drive. You can find the usable space of a specific drive and the minimum number of disks or disk partitions required for each RAID policy in *Hardware Universe*. You can also use the `storage aggregate show-spare-disks` command to find the usable space of a specific disk.

In addition to the number of disks or disk partitions necessary to create your RAID group and meet your capacity requirements, you should also have the minimum number of hot spare disks or hot spare disk partitions recommended for your aggregate:

- For all flash aggregates, you should have a minimum of one hot spare disk or disk partition.
- For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions.


```
-----
1.0.1          BSAS  7200 block      753.8GB      0B  828.0GB
1.0.10         BSAS  7200 block           0B  73.89GB  828.0GB
-----
```

You start the disk replacement job:

```
cl1:> storage disk replace -disk 1.0.1 -replacement 1.0.10 -action start
```

While you are waiting for the replacement operation to finish, you display the progress of the operation:

```
cl1:> storage aggregate show-status -aggregate aggr0_1

Owner Node: c1-01
Aggregate: aggr0_1 (online, raid_dp) (block checksums)
Plex: /aggr0_1/plex0 (online, normal, active, pool0)
RAID Group /aggr0_1/plex0/rg0 (normal, block checksums)

-----
Position Disk                Pool Type   RPM      Usable Physical
-----
shared  1.0.1                  0  BSAS    7200    73.89GB  828.0GB (replacing, copy in
progress)
shared  1.0.10                  0  BSAS    7200    73.89GB  828.0GB (copy 63% completed)
shared  1.0.0                    0  BSAS    7200    73.89GB  828.0GB (normal)
shared  1.0.11                   0  BSAS    7200    73.89GB  828.0GB (normal)
shared  1.0.6                      0  BSAS    7200    73.89GB  828.0GB (normal)
shared  1.0.5                      0  BSAS    7200    73.89GB  828.0GB (normal)
-----
```

After the replacement operation is complete, you confirm that you have a full spare disk:

```
ie2220:> storage aggregate show-spare-disks -original-owner c1-01

Original Owner: c1-01
Pool0
  Shared HDD Spares

-----
Disk                Type   RPM Checksum      Local   Local
-----
1.0.1              BSAS  7200 block      753.8GB 73.89GB 828.0GB
-----
```

Determining drive and RAID group information for an aggregate

Some aggregate administration tasks require that you know what types of drives compose the aggregate, their size, checksum, and status, whether they are shared with other aggregates, and the size and composition of the RAID groups.

Step

1. Show the drives for the aggregate, by RAID group:

```
storage aggregate show-status aggr_name
```

The drives are displayed for each RAID group in the aggregate.

You can see the RAID type of the drive (data, parity, dparity) in the `Position` column. If the `Position` column displays `shared`, then the drive is shared: if it is an HDD, it is a partitioned disk; if it is an SSD, it is part of a storage pool.

Example: A Flash Pool aggregate using an SSD storage pool and data partitions

```
cluster1:> storage aggregate show-status nodeA_fp_1

Owner Node: cluster1-a
Aggregate: nodeA_fp_1 (online, mixed_raid_type, hybrid) (block checksums)
Plex: /nodeA_fp_1/plex0 (online, normal, active, pool0)
RAID Group /nodeA_fp_1/plex0/rg0 (normal, block checksums, raid_dp)

-----
Position Disk                Pool Type   RPM      Usable Physical
-----
shared  2.0.1                      0  SAS    10000   472.9GB 547.1GB (normal)
-----
```

```

shared 2.0.3          0 SAS 10000 472.9GB 547.1GB (normal)
shared 2.0.5          0 SAS 10000 472.9GB 547.1GB (normal)
shared 2.0.7          0 SAS 10000 472.9GB 547.1GB (normal)
shared 2.0.9          0 SAS 10000 472.9GB 547.1GB (normal)
shared 2.0.11         0 SAS 10000 472.9GB 547.1GB (normal)

RAID Group /nodeA_flashpool_1/plex0/rg1 (normal, block checksums, raid4) (Storage
Pool: SmallSP)
-----
Position Disk          Pool Type      RPM      Usable Physical
-----
shared 2.0.13          0 SSD -      186.2GB 745.2GB (normal)
shared 2.0.12          0 SSD -      186.2GB 745.2GB (normal)
8 entries were displayed.

```

Relocating aggregate ownership within an HA pair

You can change the ownership of aggregates among the nodes in an HA pair without interrupting service from the aggregates.

Both nodes in an HA pair are physically connected to each other's disks or array LUNs. Each disk or array LUN is owned by one of the nodes. Although ownership of disks temporarily changes when a takeover occurs, the aggregate relocation operations either permanently (for example, if done for load balancing) or temporarily (for example, if done as part of takeover) change the ownership of all disks or array LUNs within an aggregate from one node to the other. The ownership changes without any data-copy processes or physical movement of the disks or array LUNs.

Relocating aggregate ownership

You can change the ownership of an aggregate only between the nodes within an HA pair.

About this task

- Because volume count limits are validated programmatically during aggregate relocation operations, it is not necessary to check for this manually.
If the volume count exceeds the supported limit, the aggregate relocation operation fails with a relevant error message.
- You should not initiate aggregate relocation when system-level operations are in progress on either the source or the destination node; likewise, you should not start these operations during the aggregate relocation.

These operations can include the following:

- Takeover
- Giveback
- Shutdown
- Another aggregate relocation operation
- Disk ownership changes
- Aggregate or volume configuration operations
- Storage controller replacement
- ONTAP upgrade
- ONTAP revert
- If you have a MetroCluster configuration, you should not initiate aggregate relocation while disaster recovery operations (*switchover*, *healing*, or *switchback*) are in progress.

- If you have a MetroCluster configuration and initiate aggregate relocation on a switched-over aggregate, the operation might fail because it exceeds the DR partner's volume limit count.
- You should not initiate aggregate relocation on aggregates that are corrupt or undergoing maintenance.
- If the source node is used by an Infinite Volume with SnapDiff enabled, you must perform additional steps before initiating the aggregate relocation and then perform the relocation in a specific manner.
You must ensure that the destination node has a namespace mirror constituent and make decisions about relocating aggregates that include namespace constituents.
[Infinite volumes management](#)
- Before initiating the aggregate relocation, you should save any core dumps on the source and destination nodes.

Steps

1. View the aggregates on the node to confirm which aggregates to move and ensure they are online and in good condition:

```
storage aggregate show -node source-node
```

Example

The following command shows six aggregates on the four nodes in the cluster. All aggregates are online. Node1 and Node3 form an HA pair and Node2 and Node4 form an HA pair.

```
cluster::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes  RAID Status
-----
aggr_0         239.0GB   11.13GB   95% online    1 node1  raid_dp, normal
aggr_1         239.0GB   11.13GB   95% online    1 node1  raid_dp, normal
aggr_2         239.0GB   11.13GB   95% online    1 node2  raid_dp, normal
aggr_3         239.0GB   11.13GB   95% online    1 node2  raid_dp, normal
aggr_4         239.0GB   238.9GB    0% online    5 node3  raid_dp, normal
aggr_5         239.0GB   239.0GB    0% online    4 node4  raid_dp, normal
6 entries were displayed.
```

2. Issue the command to start the aggregate relocation:

```
storage aggregate relocation start -aggregate-list aggregate-1,
aggregate-2... -node source-node -destination destination-node
```

The following command moves the aggregates aggr_1 and aggr_2 from Node1 to Node3. Node3 is Node1's HA partner. The aggregates can be moved only within the HA pair.

```
cluster::> storage aggregate relocation start -aggregate-list aggr_1,
aggr_2 -node node1 -destination node3
Run the storage aggregate relocation show command to check relocation
status.
node1::storage aggregate>
```

3. Monitor the progress of the aggregate relocation with the `storage aggregate relocation show` command:

```
storage aggregate relocation show -node source-node
```

Example

The following command shows the progress of the aggregates that are being moved to Node3:

```
cluster::> storage aggregate relocation show -node node1
Source Aggregate   Destination   Relocation Status
-----
node1
    aggr_1         node3        In progress, module: waf1
    aggr_2         node3        Not attempted yet
2 entries were displayed.
node1::storage aggregate>
```

When the relocation is complete, the output of this command shows each aggregate with a relocation status of Done.

Commands for aggregate relocation

There are specific ONTAP commands for relocating aggregate ownership within an HA pair.

If you want to...	Use this command...
Start the aggregate relocation process	<code>storage aggregate relocation start</code>
Monitor the aggregate relocation process	<code>storage aggregate relocation show</code>

Related information

[ONTAP 9 commands](#)

Assigning aggregates to SVMs

If you assign one or more aggregates to a storage virtual machine (SVM, formerly known as Vserver), then you can use only those aggregates to contain volumes for that SVM. Assigning aggregates to your SVMs is particularly important in a multi-tenancy environment or when you use Infinite Volumes.

Before you begin

The SVM and the aggregates you want to assign to that SVM must already exist.

About this task

Assigning aggregates to your SVMs helps you keep your SVMs isolated from each other; this is especially important in a multi-tenancy environment. If you use Infinite Volumes, or plan to use them in the future, you must assign aggregates to your SVMs to keep your Infinite Volumes from impacting each other and any FlexVol volumes in your cluster.

Steps

1. Check the list of aggregates already assigned to the SVM:

```
vserver show -fields aggr-list
```

The aggregates currently assigned to the SVM are displayed. If there are no aggregates assigned, “-” is displayed.

2. Add or remove assigned aggregates, depending on your requirements:

Volume Footprints	101.0MB	0%
Aggregate Metadata	300KB	0%
Snapshot Reserve	5.98GB	5%
Total Used	6.07GB	5%
Total Physical Used	34.82KB	0%

Determining which volumes reside on an aggregate

You might need to determine which FlexVol volumes or Infinite Volume constituents reside on an aggregate before performing operations on the aggregate, such as relocating it or taking it offline.

About this task

Infinite Volume constituents are somewhat similar to FlexVol volumes, but you usually do not manage them directly. For more information about Infinite Volumes and constituents, see the *ONTAP Infinite Volumes Management Guide*.

Step

1. Enter the appropriate command, depending on whether your system has Infinite Volumes:

If your system...	Then use this command...
Does not have Infinite Volumes	<code>volume show -aggregate <i>aggregate_name</i></code>
Has Infinite Volumes	<code>volume show -is-constituent * -aggregate <i>aggregate_name</i></code>

All volumes (and, if you have Infinite Volumes, constituents) that reside on the specified aggregate are displayed.

Related information

[Infinite volumes management](#)

How you can determine and control a volume's space usage in the aggregate

You can determine which FlexVol volumes are using the most space in the aggregate and specifically which features within the volume. The `volume show-footprint` command provides information about a volume's footprint, or its space usage within the containing aggregate.

The `volume show-footprint` command shows details about the space usage of each volume in an aggregate, including offline volumes. This command bridges the gap between the output of the `volume show-space` and `aggregate show-space` commands. All percentages are calculated as a percent of aggregate size.

The following example shows the `volume show-footprint` command output for a volume called `testvol`:

```
cluster1::> volume show-footprint testvol

Vserver : thevs
Volume  : testvol

Feature                               Used      Used%
-----
```


Volume Data Footprint	120.6MB	4%
Volume Guarantee	1.88GB	71%
Flexible Volume Metadata	11.38MB	0%
Delayed Frees	1.36MB	0%
Total Footprint	2.01GB	76%

The following table explains some of the key rows of the output of the `volume show-footprint` command and what you can do to try to decrease space usage by that feature:

Row/feature name	Description/contents of row	Some ways to decrease
Volume Data Footprint	The total amount of space used in the containing aggregate by a volume's data in the active file system and the space used by the volume's Snapshot copies. This row does not include reserved space.	<ul style="list-style-type: none"> Deleting data from the volume. Deleting Snapshot copies from the volume.
Volume Guarantee	The amount of space reserved by the volume in the aggregate for future writes. The amount of space reserved depends on the guarantee type of the volume.	Changing the type of guarantee for the volume to none .
Flexible Volume Metadata	The total amount of space used in the aggregate by the volume's metadata files.	No direct method to control.
Delayed Frees	Blocks that ONTAP used for performance and cannot be immediately freed. For SnapMirror destinations, this row has a value of 0 and is not displayed.	No direct method to control.
File Operation Metadata	The total amount of space reserved for file operation metadata.	No direct method to control.
Total Footprint	The total amount of space that the volume uses in the aggregate. It is the sum of all of the rows.	Any of the methods used to decrease space used by a volume.

Related information

[NetApp Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment](#)

Methods to create space in an aggregate

If an aggregate runs out of free space, various problems can result that range from loss of data to disabling a volume's guarantee. There are multiple ways to make more space in an aggregate.

All of the methods have various consequences. Prior to taking any action, you should read the relevant section in the documentation.

The following are some common ways to make space in an aggregate, in order of least to most consequences:

- Add disks to the aggregate.
- Move some volumes to another aggregate with available space.
- Shrink the size of volume-guaranteed volumes in the aggregate.
You can do this manually or with the `autoshrink` option of the `autosize` capability.
- Change volume guarantee types to **none** on volumes that are using large amounts of space (large volume-guaranteed volumes with large reserved files) so that the volumes take up less space in the aggregate.
A volume with a guarantee type of **none** has a smaller footprint in the aggregate than a volume with a guarantee type of **volume**.
- Delete unneeded volume Snapshot copies if the volume's guarantee type is **none**.
- Delete unneeded volumes.
- Enable space-saving features, such as deduplication or compression.
- (Temporarily) disable features that are using a large amount of metadata .

Commands for managing aggregates

You use the `storage aggregate` command to manage your aggregates.

If you want to...	Use this command...
Display the size of the cache for all Flash Pool aggregates	<code>storage aggregate show -fields hybrid-cache-size-total -hybrid-cache-size-total >0</code>
Display disk information and status for an aggregate	<code>storage aggregate show-status</code>
Display spare disks by node	<code>storage aggregate show-spare-disks</code>
Display the root aggregates in the cluster	<code>storage aggregate show -has-mroot true</code>
Display basic information and status for aggregates	<code>storage aggregate show</code>
Display the type of storage used in an aggregate	<code>storage aggregate show -fields storage-type</code>
Bring an aggregate online	<code>storage aggregate online</code>
Delete an aggregate	<code>storage aggregate delete</code>
Put an aggregate into the restricted state	<code>storage aggregate restrict</code>
Rename an aggregate	<code>storage aggregate rename</code>
Take an aggregate offline	<code>storage aggregate offline</code>
Change the RAID type for an aggregate	<code>storage aggregate modify -raidtype</code>

Related information

[ONTAP 9 commands](#)

Working with Flash Pool aggregates

You configure a Flash Pool aggregate by adding SSD RAID groups to an existing HDD aggregate. The SSD RAID groups can be composed of physical SSDs or allocation units from SSD storage pools.

Flash Pool caching policies and SSD partitioning

Caching policies for the volumes in a Flash Pool aggregate let you deploy flash as high performance cache for your working data set while using lower-cost HDDs for less frequently accessed data. If you are providing cache to two or more Flash Pool aggregates, you should use Flash Pool SSD partitioning to share SSDs across the aggregates in the Flash Pool.

How Flash Pool aggregate caching policies work

Caching policies are applied to volumes that reside in Flash Pool aggregates. You should understand how caching policies work before changing them.

In most cases, the default caching policy of `auto` is the best caching policy to use. The caching policy should be changed only if a different policy provides better performance for your workload. Configuring the wrong caching policy can severely degrade volume performance; the performance degradation could increase gradually over time.

Caching policies combine a read caching policy and a write caching policy. The policy name concatenates the names of the read caching policy and the write caching policy, separated by a hyphen. If there is no hyphen in the policy name, the write caching policy is “none”, except for the `auto` policy.

Read caching policies optimize for future read performance by placing a copy of the data in the cache in addition to the stored data on HDDs. For read caching policies that insert data into the cache for write operations, the cache operates as a *write-through* cache.

Data inserted into the cache by using the write caching policy exists only in cache; there is no copy in HDDs. Flash Pool cache is RAID protected. Enabling write caching makes data from write operations available for reads from cache immediately, while deferring writing the data to HDDs until it ages out of the cache.

You can change the caching policy for a volume that resides on a Flash Pool aggregate by using the `-caching-policy` parameter with the `volume create` command. When you create a volume on a Flash Pool aggregate, by default, the `auto` caching policy is assigned to the volume.

If you move a volume from a Flash Pool aggregate to a single-tier aggregate, it loses its caching policy; if you later move it back to a Flash Pool aggregate, it is assigned the default caching policy of `auto`. If you move a volume between two Flash Pool aggregates, the caching policy is preserved.

How Flash Pool SSD partitioning works for Flash Pool aggregates using storage pools

If you are providing cache to two or more Flash Pool aggregates, you should use Flash Pool Solid-State Drive (SSD) partitioning. Flash Pool SSD partitioning allows SSDs to be shared by all the aggregates using the Flash Pool. This spreads the cost of parity over multiple aggregates, increases SSD cache allocation flexibility, and maximizes SSD performance.

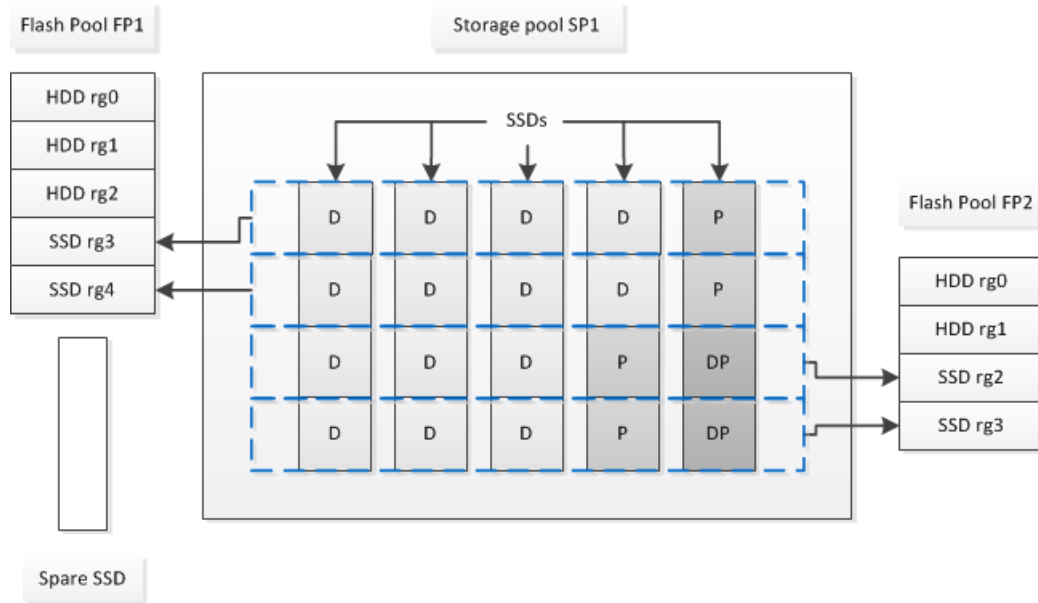
For an SSD to be used in a Flash Pool aggregate, the SSD must be placed in a storage pool. You cannot use SSDs that have been partitioned for root-data partitioning in a storage pool. After the SSD is placed in the storage pool, the SSD can no longer be managed as a stand-alone disk and cannot be

removed from the storage pool unless you destroy the aggregates associated with the Flash Pool and you destroy the storage pool.

SSD storage pools are divided into four equal allocation units. SSDs added to the storage pool are divided into four partitions and one partition is assigned to each of the four allocation units. The SSDs in the storage pool must be owned by the same HA pair. By default, two allocation units are assigned to each node in the HA pair. Allocation units must be owned by the node that owns the aggregate it is serving. If more Flash cache is required for aggregates on one of the nodes, the default number of allocation units can be shifted to decrease the number on one node and increase the number on the partner node.

You can use only one spare SSD for a storage pool. If the storage pool provides allocation units to Flash Pool aggregates owned by both nodes in the HA pair, then the spare SSD can be owned by either node. However, if the storage pool provides allocation units only to Flash Pool aggregates owned by one of the nodes in the HA pair, then the SSD spare must be owned by that same node.

The following illustration is an example of Flash Pool SSD partitioning. The SSD storage pool provides cache to two Flash Pool aggregates:



Storage pool SP1 is composed of five SSDs and a hot spare SSD. Two of the storage pool's allocation units are allocated to Flash Pool FP1, and two are allocated to Flash Pool FP2. FP1 has a cache RAID type of RAID4. Therefore, the allocation units provided to FP1 contain only one partition designated for parity. FP2 has a cache RAID type of RAID-DP. Therefore, the allocation units provided to FP2 include a parity partition and a double-parity partition.

In this example, two allocation units are allocated to each Flash Pool aggregate. However, if one Flash Pool aggregate required a larger cache, you could allocate three of the allocation units to that Flash Pool aggregate, and only one to the other.

Determining Flash Pool candidacy and optimal cache size

Before converting an existing aggregate to a Flash Pool aggregate, you can determine whether the aggregate is I/O bound, and what would be the best Flash Pool cache size for your workload and budget. You can also check whether the cache of an existing Flash Pool aggregate is sized correctly.

Before you begin

You should know approximately when the aggregate you are analyzing experiences its peak load.

Steps

1. Enter advanced mode:

```
set advanced
```

2. If you need to determine whether an existing aggregate would be a good candidate for conversion to a Flash Pool aggregate, determine how busy the disks in the aggregate are during a period of peak load, and how that is affecting latency:

```
statistics show-periodic -object disk:raid_group -instance
raid_group_name -counter disk_busy|user_read_latency -interval 1 -
iterations 60
```

You can decide whether reducing latency by adding Flash Pool cache makes sense for this aggregate.

Example

The following command shows the statistics for the first RAID group of the aggregate “aggr1”:

```
statistics show-periodic -object disk:raid_group -instance /aggr1/
plex0/rg0 -counter disk_busy|user_read_latency -interval 1 -iterations
60
```

3. Start Automated Workload Analyzer (AWA):

```
system node run -node node_name waf1 awa start aggr_name
```

AWA begins collecting workload data for the volumes associated with the specified aggregate.

4. Exit advanced mode:

```
set admin
```

Allow AWA to run until one or more intervals of peak load have occurred. AWA collects workload statistics for the volumes associated with the specified aggregate, and analyzes data for up to one rolling week in duration. Running AWA for more than one week will report only on data collected from the most recent week. Cache size estimates are based on the highest loads seen during the data collection period; the load does not need to be high for the entire data collection period.

5. Enter advanced mode:

```
set advanced
```

6. Display the workload analysis:

```
system node run -node node_name waf1 awa print
```

You can use the `-t` option to show information about the volumes in the aggregate that are the best candidates for being on a Flash Pool aggregate.

AWA displays the workload statistics and optimal Flash Pool cache size.

7. Stop AWA:

```
system node run -node node_name waf1 awa stop
```

All workload data is flushed and is no longer available for analysis.

8. Exit advanced mode:

```
set admin
```

Example

In the following example, AWA was run on aggregate “aggr1”. Here is the output of the `awa print` command after AWA had been running for about 3 days (442 10-minute intervals):

```
### FP AWA Stats ###
                                Host lada66a                                Memory 93788 MB
                                ONTAP Version NetApp Release
R8_3_1x_awa_2809198_1504220853: Tue Apr 21 18:35:12 PDT 2015

Basic Information
    Aggregate lada66a_aggr1
    Current-time Thu Apr 23 16:42:17 PDT 2015
    Start-time Thu Apr 23 08:03:51 PDT 2015
    Total runtime (sec) 31103
    Interval length (sec) 600
    Total intervals 54
    In-core Intervals 1024

Summary of the past 20 intervals
                                max
                                -----
    Read Throughput (MB/s): 181.772
    Write Throughput (MB/s): 550.611
    Cacheable Read (%): 12
    Cacheable Write (%): 30
    Max Projected Cache Size (GiB): 787.077

Summary Cache Hit Rate vs. Cache Size
Referenced Cache Size (GiB): 787.077
Referenced Interval: ID 53 starting at Thu Apr 23 16:33:07 PDT 2015
    Size      20%      40%      60%      80%
100%
  Read Hit (%)          9         20         28         32
35
  Write Hit (%)        17         21         23         25         30
```

The results provide the following pieces of information:

- **Read Throughput and Write Throughput**
The throughput measurements can help you identify an aggregate that is receiving a higher amount of traffic. Note that these numbers do not indicate whether that aggregate is I/O bound.
- **Max Projected Cache Size**
The size at which the SSD cache would hold every eligible data block that was requested from disk during the AWA run. Note that this does not guarantee a hit for all future I/O operations, because they might request data that is not in the cache. However, if the workload during the AWA run was a typical one, and if your budget allows for it, this would be an ideal size for your Flash Pool cache.
- **Projected Read Offload and Projected Write Offload**
The approximate percentages of read and write operations that would have been handled by a Flash Pool cache of the optimal size rather than going to disk (projected cache hit rate).

Note that this number is related to the performance increase you would see by converting the aggregate to a Flash Pool aggregate, but not an exact prediction.

- **Summary Cache Hit Rate vs. Cache Size**

This table can help you predict the performance impact of decreasing the size of the SSD cache from `Max Projected Cache Size`. These values are highly impacted by your workload. Depending on whether data that was aged out of the cache was ever accessed again, the impact of decreasing the size of the cache might be large or almost nonexistent. You can use this table to find the right balance between cost and performance for your workload and budget.

Creating a Flash Pool aggregate using physical SSDs

You create a Flash Pool aggregate by enabling the feature on an existing aggregate composed of HDD RAID groups, and then adding one or more SSD RAID groups to that aggregate. This results in two sets of RAID groups for that aggregate: SSD RAID groups (the SSD cache) and HDD RAID groups.

Before you begin

- You must have identified a valid aggregate composed of HDDs to convert to a Flash Pool aggregate.
- You must have determined write-caching eligibility of the volumes associated with the aggregate, and completed any required steps to resolve eligibility issues.
- You must have determined the SSDs you will be adding, and these SSDs must be owned by the node on which you are creating the Flash Pool aggregate.
- You must have determined the checksum types of both the SSDs you are adding and the HDDs already in the aggregate.
- You must have determined the number of SSDs you are adding and the optimal RAID group size for the SSD RAID groups.
Using fewer RAID groups in the SSD cache reduces the number of parity disks required, but larger RAID groups require RAID-DP.
- You must have determined the RAID level you want to use for the SSD cache.
- You must have determined the maximum cache size for your system and determined that adding SSD cache to your aggregate will not cause you to exceed it.
- You must have familiarized yourself with the configuration requirements for Flash Pool aggregates.

About this task

After you add an SSD cache to an aggregate to create a Flash Pool aggregate, you cannot remove the SSD cache to convert the aggregate back to its original configuration.

By default, the RAID level of the SSD cache is the same as the RAID level of the HDD RAID groups. You can override this default selection by specifying the `raidtype` option when you add the first SSD RAID groups.

Steps

1. Mark the aggregate as eligible to become a Flash Pool aggregate:

```
storage aggregate modify -aggregate aggr_name -hybrid-enabled true
```

If this step does not succeed, determine write-caching eligibility for the target aggregate.

2. Add the SSDs to the aggregate by using the `storage aggregate add` command.

You can specify the SSDs by ID or by using the `diskcount` and `disktype` parameters.

If the HDDs and the SSDs do not have the same checksum type, or if the aggregate is a mixed-checksum aggregate, then you must use the `checksumstyle` parameter to specify the checksum type of the disks you are adding to the aggregate.

You can specify a different RAID type for the SSD cache by using the `raidtype` parameter.

If you want the cache RAID group size to be different from the default for the RAID type you are using, you should change it now, by using the `-cache-raid-group-size` parameter.

Creating a Flash Pool aggregate using SSD storage pools

You can configure a Flash Pool aggregate by adding one or more allocation units from an SSD storage pool to an existing HDD aggregate.

Related concepts

[How Flash Pool SSD partitioning works for Flash Pool aggregates using storage pools](#) on page 27

Determining whether a Flash Pool aggregate is using an SSD storage pool

You manage Flash Pool aggregates differently when they use SSD storage pools to provide their cache than when they use discrete SSDs.

Step

1. Display the aggregate's drives by RAID group:

```
storage aggregate show-status aggr_name
```

If the aggregate is using one or more SSD storage pools, the value for the `Position` column for the SSD RAID groups is displayed as `Shared`, and the name of the storage pool is displayed next to the RAID group name.

Creating an SSD storage pool

You can create solid state drive (SSD) storage pools to provide SSD cache for two to four Flash Pool aggregates.

About this task

- You must supply a disk list when creating or adding disks to a storage pool. Storage pools do not support a `diskcount` parameter.
- The SSDs used in the storage pool should be the same size.

Steps

1. Determine the names of the available spare SSDs:

```
storage aggregate show-spare-disks -disk-type SSD
```

The SSDs used in a storage pool can be owned by either node of an HA pair.

2. Create the storage pool:

```
storage pool create -storage-pool sp_name -disk-list disk1,disk2,...
```


3. Optional: Verify the newly created storage pool:

```
storage pool show -storage-pool sp_name
```

Result

After the SSDs are placed into the storage pool, they no longer appear as spares on the cluster, even though the storage provided by the storage pool has not yet been allocated to any Flash Pool caches. You cannot add SSDs to a RAID group as discrete drives; their storage can be provisioned only by using the allocation units of the storage pool to which they belong.

Creating a Flash Pool aggregate using SSD storage pool allocation units

You can configure a Flash Pool aggregate by adding one or more allocation units from an SSD storage pool to an existing HDD aggregate.

Before you begin

- You must have identified a valid aggregate composed of HDDs to convert to a Flash Pool aggregate.
- You must have determined write-caching eligibility of the volumes associated with the aggregate, and completed any required steps to resolve eligibility issues.
- You must have created an SSD storage pool to provide the SSD cache to this Flash Pool aggregate.
Any allocation unit from the storage pool that you want to use must be owned by the same node that owns the Flash Pool aggregate.
- You must have determined how much cache you want to add to the aggregate.
You add cache to the aggregate by allocation units. You can increase the size of the allocation units later by adding SSDs to the storage pool if there is room.
- You must have determined the RAID type you want to use for the SSD cache.
After you add a cache to the aggregate from SSD storage pools, you cannot change the RAID type of the cache RAID groups.
- You must have determined the maximum cache size for your system and determined that adding SSD cache to your aggregate will not cause you to exceed it.
You can see the amount of cache that will be added to the total cache size by using the `storage pool show` command.
- You must have familiarized yourself with the configuration requirements for Flash Pool aggregates.

About this task

If you want the RAID type of the cache to differ from that of the HDD RAID groups, you must specify the cache RAID type when you add the SSD capacity. After you add the SSD capacity to the aggregate, you can no longer change the RAID type of the cache.

After you add an SSD cache to an aggregate to create a Flash Pool aggregate, you cannot remove the SSD cache to convert the aggregate back to its original configuration.

Steps

1. Mark the aggregate as eligible to become a Flash Pool aggregate:

```
storage aggregate modify -aggregate aggr_name -hybrid-enabled true
```

If this step does not succeed, determine write-caching eligibility for the target aggregate.

2. Show the available SSD storage pool allocation units:

```
storage pool show-available-capacity
```

3. Add the SSD capacity to the aggregate:

```
storage aggregate add aggr_name -storage-pool sp_name -allocation-units
number_of_units
```

If you want the RAID type of the cache to be different from that of the HDD RAID groups, you must change it when you enter this command by using the `raidtype` parameter.

You do not need to specify a new RAID group; ONTAP automatically puts the SSD cache into separate RAID groups from the HDD RAID groups.

You cannot set the RAID group size of the cache; it is determined by the number of SSDs in the storage pool.

The cache is added to the aggregate and the aggregate is now a Flash Pool aggregate. Each allocation unit added to the aggregate becomes its own RAID group.

4. Optional: Confirm the presence and size of the SSD cache:

```
storage aggregate show aggr_name
```

The size of the cache is listed under `Total Hybrid Cache Size`.

Related concepts

[How Flash Pool SSD partitioning works for Flash Pool aggregates using storage pools](#) on page 27

Related information

[NetApp Technical Report 4070: Flash Pool Design and Implementation Guide](#)

Determining the impact to cache size of adding SSDs to an SSD storage pool

If adding SSDs to a storage pool causes your platform model's cache limit to be exceeded, ONTAP does not allocate the newly added capacity to any Flash Pool aggregates. This can result in some or all of the newly added capacity being unavailable for use.

About this task

When you add SSDs to an SSD storage pool that has allocation units already allocated to Flash Pool aggregates, you increase the cache size of each of those aggregates and the total cache on the system. If none of the storage pool's allocation units have been allocated, adding SSDs to that storage pool does not affect the SSD cache size until one or more allocation units are allocated to a cache.

Steps

1. Determine the usable size of the SSDs you are adding to the storage pool:


```
storage disk show disk_name -fields usable-size
```
2. Determine how many allocation units remain unallocated for the storage pool:


```
storage pool show-available-capacity sp_name
```

All unallocated allocation units in the storage pool are displayed.
3. Calculate the amount of cache that will be added by applying the following formula:

$$(4 - \text{number of unallocated allocation units}) \times 25\% \times \text{usable size} \times \text{number of SSDs}$$

Adding SSDs to an SSD storage pool

When you add solid state drives (SSDs) to an SSD storage pool, you increase the storage pool's physical and usable sizes and allocation unit size. The larger allocation unit size also affects allocation units that have already been allocated to Flash Pool aggregates.

Before you begin

You must have determined that this operation will not cause you to exceed the cache limit for your HA pair. ONTAP does not prevent you from exceeding the cache limit when you add SSDs to an SSD storage pool, and doing so can render the newly added storage capacity unavailable for use.

About this task

When you add SSDs to an existing SSD storage pool, the SSDs must be owned by one node or the other of the same HA pair that already owned the existing SSDs in the storage pool. You can add SSDs that are owned by either node of the HA pair.

The SSD you add to the storage pool must be the same size as disk currently used in the storage pool.

Steps

1. Optional: View the current allocation unit size and available storage for the storage pool:

```
storage pool show -instance sp_name
```

2. Find available SSDs:

```
storage disk show -container-type spare -type SSD
```

3. Add the SSDs to the storage pool:

```
storage pool add -storage-pool sp_name -disk-list disk1,disk2...
```

The system displays which Flash Pool aggregates will have their size increased by this operation and by how much, and prompts you to confirm the operation.

Commands for managing SSD storage pools

ONTAP provides the `storage pool` command for managing SSD storage pools.

If you want to...	Use this command...
Display how much storage a storage pool is providing to which aggregates	<code>storage pool show-aggregate</code>
Display how much cache would be added to the overall cache capacity for both RAID types (allocation unit data size)	<code>storage pool show -instance</code>
Display the disks in a storage pool	<code>storage pool show-disks</code>
Display the unallocated allocation units for a storage pool	<code>storage pool show-available-capacity</code>
Change the ownership of one or more allocation units of a storage pool from one HA partner to the other	<code>storage pool reassign</code>

Related information

[ONTAP 9 commands](#)

Determining whether to modify the caching policy of Flash Pool aggregates

You can assign cache-retention policies to volumes in Flash Pool aggregates to determine how long the volume data remains in the Flash Pool cache. However, in some cases changing the cache-retention policy might not impact the amount of time the volume's data remains in the cache.

About this task

If your data meets any of the following conditions, changing your cache-retention policy might not have an impact:

- Your workload is sequential.
- Your workload does not reread the random blocks cached in the solid state drives (SSDs).
- The cache size of the volume is too small.

The following steps check for these conditions. The task must be done in advanced privilege mode.

Steps

1. View the workload volume:

```
statistics start -object workload_volume
```

2. Determine the workload pattern of the volume:

```
statistics show -object workload_volume -instance volume-workload -
counters sequential_reads
```

3. Determine the hit rate of the volume:

```
statistics show -object waf1_hya_vvol -instance volume -counter
read_ops_replaced_ppercent|wc_write_blks_overwritten_percent
```

4. Determine the Cacheable Read and Project Cache Alloc of the volume:

```
system node run -node node_name waf1 awa start aggr_name
```

5. Display the AWA summary:

```
system node run -node node_name waf1 awa print aggr_name
```

6. Compare the volume's hit rate to the Cacheable Read.

If the hit rate of the volume is greater than the Cacheable Read, then your workload does not reread random blocks cached in the SSDs.

7. Compare the volume's current cache size to the Project Cache Alloc.

If the current cache size of the volume is greater than the Project Cache Alloc, then the size of your volume cache is too small.

Modifying caching policies of Flash Pool aggregates

You should modify the caching policy of a volume only if a different caching policy is expected to provide better performance. You can modify the caching policy of a volume on a Flash Pool aggregate.

Before you begin

You must determine whether you want to modify your caching policy.

About this task

In most cases, the default caching policy of **auto** is the best caching policy that you can use. The caching policy should be changed only if a different policy provides better performance for your workload. Configuring the wrong caching policy can severely degrade volume performance; the performance degradation could increase gradually over time. You should use caution when modifying caching policies. If you experience performance issues with a volume for which the caching policy has been changed, you should return the caching policy to **auto**.

Step

1. Modify the volume's caching policy:

```
volume modify -volume volume_name -caching-policy policy_name
```

Example

The following example modifies the caching policy of a volume named “vol2” to the policy **none**:

```
volume modify -volume vol2 -caching-policy none
```

Related tasks

[Determining whether to modify the caching policy of Flash Pool aggregates](#) on page 36

Setting the cache-retention policy for Flash Pool aggregates

Beginning with ONTAP 9.0, you can assign cache-retention policies to volumes in Flash Pool aggregates. Data in volumes with a high cache-retention policy remains in cache longer and data in volumes with a low cache-retention policy is removed sooner. This increases performance of your critical workloads by making high priority information accessible at a faster rate for a longer period of time.

Before you begin

You should know whether your system has any conditions that might prevent the cache-retention policy from having an impact on how long your data remains in cache.

About this task

The task must be done in advanced privilege mode.

Steps

1. Change the privilege setting to advanced:


```
set -privilege advanced
```
2. Verify the volume's cache-retention policy:

By default the cache retention policy is **normal**.
3. Set the cache-retention policy:

ONTAP Version	Command
ONTAP 9.0, 9.1	<pre>priority hybrid-cache set volume_name read- cache=read_cache_value write- cache=write_cache_value cache-retention- priority=cache_retention_policy</pre> <p>Set <i>cache_retention_policy</i> to high for data that you want to remain in cache longer. Set <i>cache_retention_policy</i> to low for data that you want to remove from cache sooner.</p>
ONTAP 9.2 or later	<pre>volume modify -volume volume_name -vserver vserver_name -caching-policy policy_name.</pre>

4. Verify that the volume's cache-retention policy is changed to the option you selected.
5. Return the privilege setting to admin:

```
set -privilege admin
```

Managing storage tiers by using FabricPool

Storing data in tiers can enhance the efficiency of your storage system. Beginning with ONTAP 9.2, you can manage storage tiers by using FabricPool. It enables you to store data in a tier, based on whether the data is frequently accessed.

Benefits of storage tiers by using FabricPool

Configuring an aggregate to use FabricPool enables you to use storage tiers. You can efficiently balance the performance and cost of your storage system, monitor and optimize the space utilization, and perform policy-based data movement between storage tiers.

- You can optimize storage performance and reduce storage cost by storing data in a tier based on whether the data is frequently accessed.
 - Frequently accessed (“hot”) data is stored in the *performance tier*.
The performance tier uses high-performance primary storage, such as an all flash (all SSD) aggregate of the storage system.
 - Infrequently accessed (“cold”) data is stored in the *cloud tier*, also known as the *capacity tier*.
The cloud tier uses an object store that is less costly and does not require high performance.
- You have the flexibility in specifying the tier in which data should be stored.
You can specify one of the supported tiering policy options at the volume level. The options enable you to efficiently move data across tiers as data becomes hot or cold.
[Types of FabricPool tiering policies](#) on page 42
- You can choose one of the supported object stores to use as the cloud tier for FabricPool.
- You can monitor the space utilization in a FabricPool-enabled aggregate.
- You can see how much data in a volume is inactive by using inactive data reporting.
- You can reduce the on-premise footprint of the storage system.
You save physical space when you use a cloud-based object store for the cloud tier.

Considerations and requirements for using FabricPool

You should familiarize yourself with a few considerations and requirements about using FabricPool.

General considerations and requirements

- You must be running ONTAP 9.2 at the minimum to use FabricPool.
- You must be running ONTAP 9.4 or later releases for the following FabricPool functionality:
 - The **auto** tiering policy
[Types of FabricPool tiering policies](#) on page 42
 - Specifying the tiering minimum cooling period
 - Inactive data reporting
 - Using Microsoft Azure Blob Storage for the cloud as the cloud tier for FabricPool
 - Using FabricPool with ONTAP Select

- You must be running ONTAP 9.5 or later releases for the following FabricPool functionality:
 - Specifying the tiering fullness threshold
 - Using IBM Cloud Object Storage as the cloud tier for FabricPool
 - NetApp Volume Encryption (NVE) of the cloud tier, enabled by default.
- FabricPool supports the following platforms and aggregate types:
 - On AFF and FAS systems, you can use only all flash (all SSD) aggregates for FabricPool. You cannot use Flash Pool aggregates (which contain both SSDs and HDDs) or HDD aggregates.
 - On Cloud Volumes ONTAP and ONTAP Select, you can use either SSD or HDD aggregates for FabricPool. However, using SSD aggregates is recommended.
- FabricPool supports using the following object stores as the cloud tier:
 - Amazon Web Services Simple Storage Service (AWS S3)
 - Microsoft Azure Blob Storage for the cloud
 - NetApp StorageGRID 10.3 or later
 - IBM Cloud Object Storage
- The object store “bucket” (container) you plan to use must have already been set up, must have at least 10 GB of storage space, and must not be renamed.
- You can attach only a single object store bucket to an aggregate. However, you can attach an object store bucket to multiple aggregates in the cluster.
- HA pairs that use FabricPool require intercluster LIFs to communicate with the object store.
- You cannot detach an object store bucket from the FabricPool configuration after it is attached.
- If you use throughput floors (QoS Min), the tiering policy on the volumes must be set to **none** before the aggregate can be attached to FabricPool. Other tiering policies prevent the aggregate from being attached to FabricPool.
- You should follow the best practice guidelines for using FabricPool in specific scenarios. [NetApp Technical Report 4598: FabricPool Best Practices in ONTAP 9](#)

Additional considerations for using AWS S3 with FabricPool

- You might need a FabricPool license.
 - Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use AWS S3 on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license. If you order FabricPool for the first time for an existing cluster, a FabricPool license with 10 TB of free capacity is available.
 - If you use Amazon S3 as the cloud tier for Cloud Volumes ONTAP, you do not need a FabricPool license.
- The LIF that ONTAP uses to connect with the AWS S3 object server must be on a 10 Gbps port.
- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Amazon S3 storage classes:

- Amazon S3 Standard
- Amazon S3 Standard - Infrequent Access (Standard - IA)

Amazon Web Services (AWS) Documentation: Amazon S3 Storage Classes

- On Cloud Volumes ONTAP, FabricPool supports tiering from General Purpose SSD (gp2) and Throughput Optimized HDD (st1) volumes of Amazon Elastic Block Store (EBS).

Additional considerations for using Azure Blob Storage with FabricPool

- You might need a FabricPool license.
Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use Azure Blob Storage on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.
If you order FabricPool for the first time for an existing cluster, a FabricPool license with 10 TB of free capacity is available.
- The LIF that ONTAP uses to connect with the Azure Blob Storage object server must be on a 10 Gbps port.
- FabricPool currently does not support Azure Stack, which is on-premises Azure services.
- At the account level in Microsoft Azure Blob Storage, FabricPool supports only hot and cool storage tiers.
FabricPool does not support blob-level tiering. It also does not support tiering to Azure's archive storage tier.

Additional considerations for using StorageGRID with FabricPool

- You need to install a CA certificate for StorageGRID, unless you explicitly disable certificate checking.
- You must not enable StorageGRID object versioning on the object store bucket.
- A FabricPool license is not required.

Additional considerations for using IBM Cloud Object Storage with FabricPool

- You might need a FabricPool license.
Newly ordered AFF systems come with 10 TB of free capacity for using FabricPool. If you need additional capacity on an AFF system, if you use IBM Cloud Object Storage on a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.
If you order FabricPool for the first time for an existing cluster, a FabricPool license with 10 TB of free capacity is available.

Functionality or features not supported by FabricPool

- Information lifecycle management (ILM) policies that are applied to object store buckets
ILM typically includes various movement and deletion policies. These policies can be disruptive to the data in the cloud tier of FabricPool. Using FabricPool with ILM policies that are configured on object stores can result in data loss.
- 7-Mode data transition using the ONTAP CLI commands or the 7-Mode Transition Tool
- FlexArray Virtualization
- Infinite Volumes
- MetroCluster configurations

You cannot use FabricPool with a MetroCluster configuration, either as a mirrored or an unmirrored aggregate.

- RAID SyncMirror
- SnapLock volumes
- SVM disaster recovery (DR)

You cannot set up SVM DR if the SVM contains any volumes that reside in a FabricPool-enabled aggregate. Likewise, if an SVM already has SVM DR enabled, you cannot create a volume on a FabricPool-enabled aggregate for that SVM or move a volume of that SVM to a FabricPool-enabled aggregate.
- Tape backup using SMTape, NDMP, or a dump backup
- The Auto Balance functionality
- Volumes using a space guarantee other than **none**

FabricPool does not support attaching a cloud tier to an aggregate that contains volumes using a space guarantee other than **none**. For example, a volume using a space guarantee of **volume** (`-space-guarantee volume`) is not supported.

Understanding FabricPool tiering policies

FabricPool tiering policies enable you to move data efficiently across tiers as data becomes hot or cold. Understanding the tiering policies helps you select the right policy that suits your storage management needs.

Types of FabricPool tiering policies

FabricPool tiering policies determine when or whether the user data blocks of a volume in FabricPool are moved to the cloud tier, based on the volume “temperature” of hot (active) or cold (inactive). The volume “temperature” increases when it is accessed frequently and decreases when it is not. Some tiering policies have an associated tiering minimum cooling period, which sets the time that user data in a volume of FabricPool must remain inactive for the data to be considered “cold” and moved to the cloud tier.

The FabricPool tiering policy is specified at the volume level. Four options are available:

- The **snapshot-only** tiering policy (the default) moves user data blocks of the volume Snapshot copies that are not associated with the active file system to the cloud tier.

The tiering minimum cooling period is 2 days. You can modify the default setting for the tiering minimum cooling period with the `-tiering-minimum-cooling-days` parameter in the advanced privilege level of the `volume create` and `volume modify` commands. Valid values are 2 to 63 days.
- The **auto** tiering policy, supported only on ONTAP 9.4 and later releases, moves cold user data blocks in both the Snapshot copies and the active file system to the cloud tier.

The default tiering minimum cooling period is 31 days and applies to the entire volume, for both the active file system and the Snapshot copies.

You can modify the default setting for the tiering minimum cooling period with the `-tiering-minimum-cooling-days` parameter in the advanced privilege level of the `volume create` and `volume modify` commands. Valid values are 2 to 63 days.
- The **backup** tiering policy enables newly transferred data of a data protection volume to start in the cloud tier.

The tiering minimum cooling period does not apply because the data starts in the cloud tier, and you cannot modify the setting.

- The **none** tiering policy keeps data of a volume in the performance tier, preventing it from being moved to the cloud tier.
The tiering minimum cooling period does not apply because the data never moves to the cloud tier, and you cannot modify the setting.

The `volume show` command output shows the tiering policy of a volume. A volume that has never been used with FabricPool shows the **snapshot-only** tiering policy in the output.

What happens when you modify the tiering policy of a volume in FabricPool

You can modify the tiering policy of a volume by performing a `volume modify` operation. You must understand how changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

- Changing the tiering policy from **snapshot-only** or **none** to **auto** causes ONTAP to send user data blocks in the active file system that are already cold to the cloud tier, even if those user data blocks were not previously eligible for the cloud tier.
- Changing the tiering policy from **auto** to **snapshot-only** or **none** does not cause active file system blocks that are already moved to the cloud tier to be moved back to the performance tier. Volume reads are needed for the data to be moved back to the performance tier.
- Any time you change the tiering policy on a volume, the tiering minimum cooling period is reset to the default value for the policy.

What happens to the tiering policy when you move a volume

- Unless you explicitly specify a different tiering policy, a volume retains its original tiering policy when it is moved in and out of a FabricPool-enabled aggregate.
However, the tiering policy takes effect only when the volume is in a FabricPool-enabled aggregate.
- The existing value of the `-tiering-minimum-cooling-days` parameter for a volume moves with the volume unless you specify a different tiering policy for the destination.
If you specify a different tiering policy, then the volume uses the default tiering minimum cooling period for that policy. This is the case whether the destination is FabricPool or not.
- You can move a volume across aggregates and at the same time modify the tiering policy.
- You should pay special attention when a `volume move` operation involves the **auto** tiering policy.

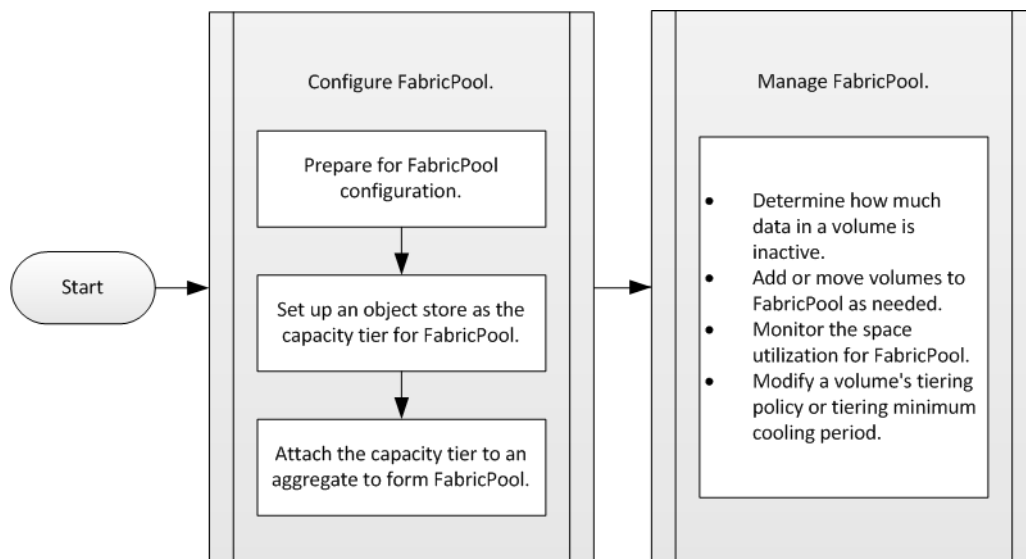
Assuming that both the source and the destination are FabricPool-enabled aggregates, the following table summarizes the outcome of a `volume move` operation that involves policy changes related to **auto**:

When you move a volume that has a tiering policy of...	And you change the tiering policy with the move to...	Then after the volume move...
backup	auto	All data is moved to the performance tier.
snapshot-only, none, or auto	auto	Data blocks are moved to the same tier of the destination as they previously were on the source.
auto,	snapshot-only	All data is moved to the performance tier.

When you move a volume that has a tiering policy of...	And you change the tiering policy with the move to...	Then after the volume move...
auto	backup	All user data is moved to the cloud tier.
auto	none	All data is kept at the performance tier.

FabricPool management workflow

You can use the FabricPool workflow diagram to help you plan the configuration and management tasks.



Configuring FabricPool

Configuring FabricPool helps you manage which storage tier (the local performance tier or the cloud tier) data should be stored based on whether the data is frequently accessed.

Steps

1. [Preparing for FabricPool configuration](#) on page 45
The preparation required for FabricPool configuration depends on the object store you use as the cloud tier.
2. [Setting up an object store as the cloud tier for FabricPool](#) on page 47
Setting up FabricPool involves specifying the configuration information of the object store (AWS S3, Azure Blob Storage for the cloud, IBM Cloud Object Storage, Alibaba Cloud Object Storage, or StorageGRID) that you plan to use as the cloud tier for FabricPool.
3. [Attaching the cloud tier to an aggregate](#) on page 50
After setting up an object store as the cloud tier, you specify the aggregate to use by attaching it to FabricPool. In ONTAP 9.5 and later, you can also attach aggregates that contain qualified FlexGroup volume constituents.

Related concepts

[Considerations and requirements for using FabricPool](#) on page 39

[FabricPool management workflow](#) on page 44

Preparing for FabricPool configuration

The preparation required for FabricPool configuration depends on the object store you use as the cloud tier.

Choices

- [Installing a FabricPool license if required for AWS S3, Azure Blob Storage, or IBM Cloud Object Storage](#) on page 45

Newly ordered AFF systems come with 10 TB of free capacity for using AWS S3, Azure Blob Storage, or IBM Cloud Object Storage with FabricPool, if you select it. If you need additional capacity on an AFF system, if you plan to use AWS S3, Azure Blob Storage, or IBM Cloud Object Storage with a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

- [Installing a CA certificate if you use StorageGRID](#) on page 46

Unless you plan to disable certificate checking for StorageGRID, you must install a StorageGRID CA certificate on the cluster so that ONTAP can authenticate with StorageGRID as the object store for FabricPool.

Installing a FabricPool license if required for AWS S3, Azure Blob Storage, or IBM Cloud Object Storage

Newly ordered AFF systems come with 10 TB of free capacity for using AWS S3, Azure Blob Storage, or IBM Cloud Object Storage with FabricPool, if you select it. If you need additional capacity on an AFF system, if you plan to use AWS S3, Azure Blob Storage, or IBM Cloud Object Storage with a non-AFF system, or if you upgrade from an existing cluster, you need a FabricPool license.

About this task

The FabricPool license is a cluster-wide license. It includes an entitled usage limit that you purchase for the AWS S3, Azure Blob Storage, or IBM Cloud Object Storage cloud tier that is associated with FabricPool in the cluster. The usage across the cluster must not exceed the capacity of the entitled usage limit. If you need to increase the usage limit of the license, you should contact your sales representative.

A FabricPool license with 10 TB of free capacity is available for first time FabricPool orders for existing clusters.

A license is not required if you use NetApp StorageGRID for the cloud tier, or if you use Amazon S3 as the cloud tier for Cloud Volumes ONTAP.

This task is supported only by uploading the license file to the cluster using OnCommand System Manager.

Steps

1. Download the NetApp License File (NLF) for the FabricPool license from the NetApp Support Site.
[NetApp Support](#)
2. Use OnCommand System Manager to upload the FabricPool license to the cluster:
 - a. Depending on the ONTAP release running on your storage system, perform the following action:

If you are running...	Perform the following action...
ONTAP 9.4 or earlier	Click Configuration > Licenses .
ONTAP 9.5 or later	Click Configuration > Cluster > Licenses .

- b. In the **Cluster Settings** pane, click **Licenses**.
- c. In the **Packages** window, click **Add**.
- d. In the **Add License Packages** dialog box, click **Choose Files** to select the NLF you downloaded, and then click **Add** to upload the file to the cluster.

Related information

[NetApp KB Article 000032437: ONTAP FabricPool Licensing Overview](#)
[NetApp Software License Search](#)

Installing a FabricPool license (video)

This video shows a quick overview of using OnCommand System Manager to install a FabricPool license for non-NetApp object stores such as AWS S3.

[NetApp video: Installing a FabricPool license](#)

Related information

[NetApp TechComm TV: FabricPool playlist](#)

Installing a CA certificate if you use StorageGRID

Unless you plan to disable certificate checking for StorageGRID, you must install a StorageGRID CA certificate on the cluster so that ONTAP can authenticate with StorageGRID as the object store for FabricPool.

About this task

ONTAP 9.4 and later releases enable you to disable certificate checking for StorageGRID.

Steps

1. Contact your StorageGRID administrator to obtain the StorageGRID system's CA certificate.
2. Use the `security certificate install` command with the `-type server-ca` parameter to install the StorageGRID CA certificate on the cluster.

The fully qualified domain name (FQDN) you enter must match the custom common name on the StorageGRID CA certificate.

Related tasks

[Setting up StorageGRID as the cloud tier](#) on page 49

If you are running ONTAP 9.2 or later, you can set up StorageGRID as the cloud tier for FabricPool.

Related information

[StorageGRID Webscale 11.1 Administrator Guide](#)

Setting up an object store as the cloud tier for FabricPool

Setting up FabricPool involves specifying the configuration information of the object store (AWS S3, Azure Blob Storage for the cloud, IBM Cloud Object Storage, Alibaba Cloud Object Storage, or StorageGRID) that you plan to use as the cloud tier for FabricPool.

Choices

- [Setting up AWS S3 as the cloud tier](#) on page 47
If you are running ONTAP 9.2 or later, you can set up AWS S3 as the cloud tier for FabricPool. If you are running ONTAP 9.5 or later, you can set up AWS Commercial Cloud Services (C2S) for FabricPool.
- [Setting up Azure Blob Storage for the cloud as the cloud tier](#) on page 48
If you are running ONTAP 9.4 or later, you can set up Azure Blob Storage for the cloud as the cloud tier for FabricPool.
- [Setting up IBM Cloud Object Storage as the cloud tier](#) on page 49
If you are running ONTAP 9.5 or later, you can set up IBM Cloud Object Storage as the cloud tier for FabricPool.
- [Setting up StorageGRID as the cloud tier](#) on page 49
If you are running ONTAP 9.2 or later, you can set up StorageGRID as the cloud tier for FabricPool.

Setting up AWS S3 as the cloud tier

If you are running ONTAP 9.2 or later, you can set up AWS S3 as the cloud tier for FabricPool. If you are running ONTAP 9.5 or later, you can set up AWS Commercial Cloud Services (C2S) for FabricPool.

Steps

1. Specify the AWS S3 configuration information by using the `storage aggregate object-store config create` command with the `-provider-type AWS_S3` parameter.
 - You use the `-auth-type CAP` parameter to obtain credentials for C2S access.
When you use the `-auth-type CAP` parameter, you must use the `-cap-url` parameter to specify the full URL to request temporary credentials for C2S access.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access AWS S3 with the provided information.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the AWS S3 object store.
 - You use the `-secret-password` parameter to specify the password (secret access key) for authenticating requests to the AWS S3 object store.
 - If the AWS S3 password is changed, you should update the corresponding password stored in ONTAP immediately.
Doing so enables ONTAP to access the data in AWS S3 without interruption.

Example

```
cluster1:~> storage aggregate object-store config create
-object-store-name my_aws_store -provider-type AWS_S3
-server s3.amazonaws.com -container-name my-aws-bucket
-access-key DXJRXHPXHYXA9X31X3JX
```

Example

```
cluster1::> storage aggregate object-store config create -object-
store-name my_c2s_store -provider-type AWS_S3 -auth-type CAP -cap-url
https://123.45.67.89/api/v1/credentials?
agency=XYZ&mission=TESTACCT&role=S3FULLACCESS -server my-c2s-s3server-
fqdn -container my-c2s-s3-bucket
```

2. Display and verify the AWS S3 configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the AWS S3 configuration information for FabricPool.

Attaching AWS S3 as a FabricPool cloud tier (video)

This video shows a quick overview of using OnCommand System Manager to attach an AWS S3 bucket to ONTAP aggregates with FabricPool.

[NetApp video: Attaching Amazon S3 as a FabricPool external capacity tier](#)

Related information

[NetApp TechComm TV: FabricPool playlist](#)

Setting up Azure Blob Storage for the cloud as the cloud tier

If you are running ONTAP 9.4 or later, you can set up Azure Blob Storage for the cloud as the cloud tier for FabricPool.

About this task

FabricPool currently does not support Azure Stack, which is on-premises Azure services.

Steps

1. Specify the Azure Blob Storage configuration information by using the `storage aggregate object-store config create` command with the `-provider-type Azure_Cloud` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access Azure Blob Storage with the provided information.
 - You use the `-azure-account` parameter to specify the Azure Blob Storage account.
 - You use the `-azure-private-key` parameter to specify the access key for authenticating requests to Azure Blob Storage.
 - If the Azure Blob Storage password is changed, you should update the corresponding password stored in ONTAP immediately.
Doing so enables ONTAP to access the data in Azure Blob Storage without interruption.

Example

```
cluster1::> storage aggregate object-store config create
-object-store-name MyAzure -provider-type Azure_Cloud
-server blob.core.windows.net -container-name myAzureContainer
-azure-account myAzureAcct -azure-private-key myAzureKey
```

2. Display and verify the Azure Blob Storage configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the Azure Blob Storage configuration information for FabricPool.

Setting up IBM Cloud Object Storage as the cloud tier

If you are running ONTAP 9.5 or later, you can set up IBM Cloud Object Storage as the cloud tier for FabricPool.

Steps

1. Specify the IBM Cloud Object Storage configuration information by using the `storage aggregate object-store config create` command with the `-provider-type IBM_COS` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access IBM Cloud Object Storage with the provided information.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the IBM Cloud Object Storage object store.
 - You use the `-secret-password` parameter to specify the password (secret access key) for authenticating requests to the IBM Cloud Object Storage object store.
 - If the IBM Cloud Object Storage password is changed, you should update the corresponding password stored in ONTAP immediately.
Doing so enables ONTAP to access the data in IBM Cloud Object Storage without interruption.

Example

```
storage aggregate object-store config create
-object-store-name MyIBM -provider-type IBM_COS
-server s3.us-east.objectstorage.softlayer.net
-container-name my-ibm-cos-bucket -access-key DXJRXHPXHYXA9X31X3JX
```

2. Display and verify the IBM Cloud Object Storage configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the IBM Cloud Object Storage configuration information for FabricPool.

Setting up StorageGRID as the cloud tier

If you are running ONTAP 9.2 or later, you can set up StorageGRID as the cloud tier for FabricPool.

Steps

1. Specify the StorageGRID configuration information by using the `storage aggregate object-store config create` command with the `-provider-type SGWS` parameter.
 - The `storage aggregate object-store config create` command fails if ONTAP cannot access StorageGRID with the provided information.
 - You use the `-access-key` parameter to specify the access key for authorizing requests to the StorageGRID object store.
 - You use the `-secret-password` parameter to specify the password (secret access key) for authenticating requests to the StorageGRID object store.
 - If the StorageGRID password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in StorageGRID without interruption.

- Setting the `-is-certificate-validation-enabled` parameter to **false** disables certificate checking for StorageGRID.

Example

```
cluster1::> storage aggregate object-store config create
-object-store-name mySGWS -provider-type SGWS -server mySGWSserver
-container-name mySGWScontainer -access-key mySGWSkey
-secret-password mySGWSpass
```

2. Display and verify the StorageGRID configuration information by using the `storage aggregate object-store config show` command.

The `storage aggregate object-store config modify` command enables you to modify the StorageGRID configuration information for FabricPool.

Attaching StorageGRID as a FabricPool cloud tier (video)

This video shows a quick overview of using OnCommand System Manager to attach a StorageGRID bucket to ONTAP aggregates with FabricPool.

[NetApp video: Attaching StorageGRID Webscale as a FabricPool external capacity tier](#)

Related information

[NetApp TechComm TV: FabricPool playlist](#)

Attaching the cloud tier to an aggregate

After setting up an object store as the cloud tier, you specify the aggregate to use by attaching it to FabricPool. In ONTAP 9.5 and later, you can also attach aggregates that contain qualified FlexGroup volume constituents.

Before you begin

When you use the ONTAP CLI to set up an aggregate for FabricPool, the aggregate must already exist.

Note: When you use OnCommand System Manager to set up an aggregate for FabricPool, you can create the aggregate and set it up to use for FabricPool at the same time.

Steps

1. Optional: To see how much data in a volume is inactive, follow the steps in *[Determining how much data in a volume is inactive by using inactive data reporting](#)* on page 51.

Seeing how much data in a volume is inactive can help you decide which aggregate to use for FabricPool.

2. Attach the object store to an aggregate by using the `storage aggregate object-store attach` command.

If the aggregate has never been used with FabricPool and it contains existing volumes, then the volumes are assigned the default **snapshot-only** tiering policy.

Example

```
cluster1::> storage aggregate object-store attach -aggregate myaggr
-object-store-name Amazon01B1
```

You can use the `allow-flexgroup true` option to attach aggregates that contain FlexGroup volume constituents.

3. Display the object store information and verify that the attached object store is available by using the `storage aggregate object-store show` command.

Example

```
cluster1::> storage aggregate object-store show
```

Aggregate	Object Store Name	Availability State
myaggr	Amazon01B1	available

Managing FabricPool

To help you with your storage tiering needs, ONTAP enables you to display how much data in a volume is inactive, add or move volumes to FabricPool, monitor the space utilization for FabricPool, or modify a volume's tiering policy or tiering minimum cooling period.

Choices

- [Determining how much data in a volume is inactive by using inactive data reporting](#) on page 51
Seeing how much data in a volume is inactive enables you to make good use of storage tiers. Information in inactive data reporting helps you decide which aggregate to use for FabricPool, whether to move a volume in to or out of FabricPool, or whether to modify the tiering policy of a volume.
- [Adding or moving volumes to FabricPool as needed](#) on page 53
You can add volumes to FabricPool by creating new volumes directly in the FabricPool-enabled aggregate or by moving existing volumes from another aggregate to the FabricPool-enabled aggregate.
- [Monitoring the space utilization for FabricPool](#) on page 55
You need to know how much data is stored in the performance and cloud tiers for FabricPool. That information helps you determine whether you need to change the tiering policy of a volume, increase the FabricPool licensed usage limit, or increase the storage space of the cloud tier.
- [Managing storage tiering by modifying a volume's tiering policy or tiering minimum cooling period](#) on page 56
You can change the tiering policy of a volume to control whether data is moved to the cloud tier when it becomes inactive (*cold*). For a volume with the `snapshot-only` or `auto` tiering policy, you can also specify the tiering minimum cooling period that user data must remain inactive before it is moved to the cloud tier.

Determining how much data in a volume is inactive by using inactive data reporting

Seeing how much data in a volume is inactive enables you to make good use of storage tiers. Information in inactive data reporting helps you decide which aggregate to use for FabricPool, whether to move a volume in to or out of FabricPool, or whether to modify the tiering policy of a volume.

Before you begin

You must be running ONTAP 9.4 or later to use the inactive data reporting functionality.

About this task

Inactive data reporting is not supported on some aggregates:

- On AFF and FAS systems, inactive data reporting is supported only on all-SSD aggregates.
- On Cloud Volumes ONTAP and ONTAP Select, inactive data reporting is supported on both SSD and HDD aggregates.
- Inactive data reporting does not support Flash Pool (hybrid) aggregates, MetroCluster aggregates, or SnapLock aggregates.

Steps

1. If the aggregate for which you want to see inactive data reporting is not used in FabricPool, enable inactive data reporting for the aggregate by using the `storage aggregate modify` command with the `-is-inactive-data-reporting-enabled true` parameter.

Example

```
cluster1::> storage aggregate modify -aggregate aggr1 -is-inactive-
data-reporting-enabled true
```

You need to explicitly enable the inactive data reporting functionality on an aggregate that is not used for FabricPool.

You cannot and do not need to enable inactive data reporting on a FabricPool-enabled aggregate because the aggregate already comes with inactive data reporting. The `-is-inactive-data-reporting-enabled` parameter does not work on FabricPool-enabled aggregates.

The `-fields is-inactive-data-reporting-enabled` parameter of the `storage aggregate show` command shows whether inactive data reporting is enabled on an aggregate.

2. To display how much data is inactive on a volume, use the `volume show` command with the `-fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-percent` parameter.

Example

```
cluster1::> volume show -fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-
percent
-----
vsrvr volume performance-tier-inactive-user-data performance-tier-inactive-user-data-percent
-----
vsim1  vol0  0B 0%
vs1    vs1rv1 0B 0%
vs1    vv1   10.34MB 0%
vs1    vv2   10.38MB 0%
4 entries were displayed.
```

- The `performance-tier-inactive-user-data` field displays how much user data stored in the aggregate is inactive.
- The `performance-tier-inactive-user-data-percent` field displays what percent of the data is inactive across the active file system and Snapshot copies.
- For an aggregate that is not used for FabricPool, it takes 31 days after you enable inactive data reporting for ONTAP to display inactive data for the entire volume, regardless of whether a tiering policy has been set or what the tiering policy is. Until the period is reached, the output shows “-” for the amount of inactive data instead of a value.
- On a volume that is part of FabricPool, what ONTAP reports as inactive depends on the tiering policy that is set on a volume.

- For the **auto** tiering policy, ONTAP reports the inactive amount of the entire volume for the period specified by the `-tiering-minimum-cooling-days` parameter in the advanced privilege level of the `volume create` or `volume modify` command. If no value is set for `-tiering-minimum-cooling-days`, then ONTAP uses the default value (31 days).
- For the **snapshot-only** tiering policy, ONTAP reports the amount of the active file system that is inactive for at least 31 days.
- For the **none** tiering policy, ONTAP reports the amount of the entire volume that is inactive for at least 31 days.
- For the **backup** tiering policy, inactive data reporting is not supported.

Adding or moving volumes to FabricPool as needed

You can add volumes to FabricPool by creating new volumes directly in the FabricPool-enabled aggregate or by moving existing volumes from another aggregate to the FabricPool-enabled aggregate.

Choices

- [Creating a volume for FabricPool](#) on page 53
When you create a volume for FabricPool, you have the option to specify a tiering policy. If no tiering policy is specified, the created volume uses the default **snapshot-only** tiering policy. For a volume with the **snapshot-only** or **auto** tiering policy, you can also specify the tiering minimum cooling period.
- [Moving a volume to FabricPool](#) on page 54
When you move a volume to FabricPool, you have the option to specify or change the tiering policy for the volume with the move.

Creating a volume for FabricPool

When you create a volume for FabricPool, you have the option to specify a tiering policy. If no tiering policy is specified, the created volume uses the default **snapshot-only** tiering policy. For a volume with the **snapshot-only** or **auto** tiering policy, you can also specify the tiering minimum cooling period.

Before you begin

- Setting a volume to use the **auto** tiering policy or specifying the tiering minimum cooling period requires ONTAP 9.4 or later.
- Using FlexGroup volumes requires ONTAP 9.5 or later.

Step

1. Create a new volume for FabricPool by using the `volume create` command.
 - The `-tiering-policy` optional parameter enables you to specify the tiering policy for the volume.
You can specify one of the following tiering policies:
 - **snapshot-only** (default)
 - **auto**
 - **backup** (deprecated)

- **none**

Types of FabricPool tiering policies on page 42

- The `-tiering-minimum-cooling-days` optional parameter in the advanced privilege level enables you to specify the tiering minimum cooling period for a volume that uses the **snapshot-only** or **auto** tiering policy.

You can specify a value between 2 and 63 for the tiering minimum cooling days.

Example of creating a volume for FabricPool

The following example creates a volume called “myvoll” in the “myFabricPool” FabricPool-enabled aggregate. The tiering policy is set to **auto** and the tiering minimum cooling period is set to 45 days:

```
cluster1::*> volume create -vserver myVS -aggregate myFabricPool
-volume myvoll -tiering-policy auto -tiering-minimum-cooling-days 45
```

Related concepts

Understanding FabricPool tiering policies on page 42

Related information

FlexGroup volumes management

Moving a volume to FabricPool

When you move a volume to FabricPool, you have the option to specify or change the tiering policy for the volume with the move.

Before you begin

You must understand how changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

What happens to the tiering policy when you move a volume on page 43

Step

1. Use the `volume move start` command to move a volume to FabricPool.

The `-tiering-policy` optional parameter enables you to specify the tiering policy for the volume.

You can specify one of the following tiering policies:

- **snapshot-only** (default)
- **auto**
- **backup**
- **none**

Types of FabricPool tiering policies on page 42

Example of moving a volume to FabricPool

The following example moves a volume named “myvol2” of the “vs1” SVM to the “dest_FabricPool” FabricPool-enabled aggregate. The volume is explicitly set to use the **none** tiering policy:

```
cluster1::> volume move start -vserver vs1 -volume myvol2
-destination-aggregate dest_FabricPool -tiering-policy none
```

Related concepts

[Understanding FabricPool tiering policies](#) on page 42

Monitoring the space utilization for FabricPool

You need to know how much data is stored in the performance and cloud tiers for FabricPool. That information helps you determine whether you need to change the tiering policy of a volume, increase the FabricPool licensed usage limit, or increase the storage space of the cloud tier.

Steps

1. Monitor the space utilization for FabricPool-enabled aggregates by using one of the following commands to display the information:

If you want to display...	Then use this command...
The used size of the cloud tier in an aggregate	<code>storage aggregate show</code> with the <code>-instance</code> parameter
Details of space utilization within an aggregate, including the object store's referenced capacity	<code>storage aggregate show-space</code> with the <code>-instance</code> parameter
Space utilization of the object stores that are attached to the aggregates, including how much license space is being used	<code>storage aggregate object-store show-space</code>
A list of volumes in an aggregate and the footprints of their data and metadata	<code>volume show-footprint</code>

In addition to using CLI commands, you can use OnCommand Unified Manager or OnCommand System Manager to monitor the space utilization.

Example

The following example shows ways of displaying space utilization and related information for FabricPool:

```
cluster1::> storage aggregate show-space -instance
Aggregate: MyFabricPool
...
Aggregate Display Name: MyFabricPool
...
Object Store Referenced Capacity Percentage: -
Object Store Metadata: -
Object Store Metadata Percent: -
```

```
Total Unreclaimed Space: -
Object Store Unreclaimed Space Percentage: -
Object Store Size: -
```

```
cluster1::> storage aggregate show -instance

Aggregate: MyFabricPool
...
Composite: true
Capacity Tier Used Size:
...
```

```
cluster1::> volume show-footprint

Vserver : vs1
Volume : rootvol

Feature                               Used      Used%
-----
Volume Footprint                      KB        %
Volume Guarantee                      MB        %
Flexible Volume Metadata              KB        %
Delayed Frees                         KB        %
Total Footprint                       MB        %

Vserver : vs1
Volume : vol

Feature                               Used      Used%
-----
Volume Footprint                      KB        %
Footprint in Performance Tier         KB        %
Footprint in Amazon01                 KB        %
Flexible Volume Metadata              MB        %
Delayed Frees                         KB        %
Total Footprint                       MB        %
...
```

- Take one of the following actions as needed:

If you want to...	Then...
Change the tiering policy of a volume	Follow the procedure in Managing storage tiering by modifying a volume's tiering policy or tiering minimum cooling period on page 56.
Increase the FabricPool licensed usage limit	Contact your NetApp or partner sales representative. NetApp Support
Increase the storage space of the cloud tier	Contact the provider of the object store that you use for the cloud tier.

Managing storage tiering by modifying a volume's tiering policy or tiering minimum cooling period

You can change the tiering policy of a volume to control whether data is moved to the cloud tier when it becomes inactive (*cold*). For a volume with the **snapshot-only** or **auto** tiering policy, you can also specify the tiering minimum cooling period that user data must remain inactive before it is moved to the cloud tier.

Before you begin

Changing a volume to the **auto** tiering policy or modifying the tiering minimum cooling period requires ONTAP 9.4 or later.

About this task

Changing the tiering policy of a volume changes only the subsequent tiering behavior for the volume. It does not retroactively move data to the cloud tier.

Changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

[What happens when you modify the tiering policy of a volume in FabricPool](#) on page 43

Steps

1. Modify the tiering policy for an existing volume by using the `volume modify` command with the `-tiering-policy` parameter:

You can specify one of the following tiering policies:

- `snapshot-only` (default)
- `auto`
- `backup`
- `none`

[Types of FabricPool tiering policies](#) on page 42

2. If the volume uses the `snapshot-only` or `auto` tiering policy and you want to modify the tiering minimum cooling period, use the `volume modify` command with the `-tiering-minimum-cooling-days` optional parameter in the advanced privilege level.

You can specify a value between 2 and 63 for the tiering minimum cooling days.

Example of modifying the tiering policy and the tiering minimum cooling period of a volume

The following example changes the tiering policy of the volume “myvol” in the SVM “vs1” to `auto` and the tiering minimum cooling period to 45 days:

```
cluster1::> volume modify -vserver vs1 -volume myvol
-tiering-policy auto -tiering-minimum-cooling-days 45
```

Related concepts

[Understanding FabricPool tiering policies](#) on page 42

Archiving volumes with FabricPool (video)

This video shows a quick overview of using OnCommand System Manager to archive a volume to a cloud tier with FabricPool.

[NetApp video: Archiving volumes with FabricPool \(backup + volume move\)](#)

Related information

[NetApp TechComm TV: FabricPool playlist](#)

Commands for managing aggregates with FabricPool

You use the `storage aggregate object-store` commands to manage object stores for FabricPool. You use the `storage aggregate` commands to manage aggregates for FabricPool. You use the `volume` commands to manage volumes for FabricPool.

If you want to...	Use this command:
Define the configuration for an object store so that ONTAP can access it	<code>storage aggregate object-store config create</code>
Modify object store configuration attributes	<code>storage aggregate object-store config modify</code>
Rename an existing object store configuration	<code>storage aggregate object-store config rename</code>
Delete the configuration of an object store	<code>storage aggregate object-store config delete</code>
Display a list of object store configurations	<code>storage aggregate object-store config show</code>
Attach an object store to an aggregate for using FabricPool	<code>storage aggregate object-store attach</code>
Attach an object store to an aggregate that contains a FlexGroup volume for using FabricPool	<code>storage aggregate object-store attach with the <code>allow-flexgroup true</code></code>
Display details of the object stores that are attached to FabricPool-enabled aggregates	<code>storage aggregate object-store show</code>
Display the aggregate fullness threshold used by the tiering scan	<code>storage aggregate object-store show with the <code>-fields tiering-fullness-threshold</code> parameter in the advanced privilege level</code>
Display space utilization of the object stores that are attached to FabricPool-enabled aggregates	<code>storage aggregate object-store show-space</code>
Enable inactive data reporting on an aggregate that is not used for FabricPool	<code>storage aggregate modify with the <code>-is-inactive-data-reporting-enabled true</code> parameter</code>
Display whether inactive data reporting is enabled on an aggregate	<code>storage aggregate show with the <code>-fields is-inactive-data-reporting-enabled</code> parameter</code>
Display information about how much user data is cold within an aggregate	<code>storage aggregate show-space with the <code>-fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-percent</code> parameter</code>

If you want to...	Use this command:
Create a volume for FabricPool, including specifying the following: <ul style="list-style-type: none"> • The tiering policy • The tiering minimum cooling period (for the snapshot-only or auto tiering policy) 	volume create <ul style="list-style-type: none"> • You use the <code>-tiering-policy</code> parameter to specify the tiering policy. • You use the <code>-tiering-minimum-cooling-days</code> parameter in the advanced privilege level to specify the tiering minimum cooling period.
Modify a volume for FabricPool, including modifying the following: <ul style="list-style-type: none"> • The tiering policy • The tiering minimum cooling period (for the snapshot-only or auto tiering policy) 	volume modify <ul style="list-style-type: none"> • You use the <code>-tiering-policy</code> parameter to specify the tiering policy. • You use the <code>-tiering-minimum-cooling-days</code> parameter in the advanced privilege level to specify the tiering minimum cooling period.
Display FabricPool information related to a volume, including the following: <ul style="list-style-type: none"> • The tiering minimum cooling period • How much user data is cold 	volume show <ul style="list-style-type: none"> • You use the <code>-fields tiering-minimum-cooling-days</code> parameter in the advanced privilege level to display the tiering minimum cooling period. • You use the <code>-fields performance-tier-inactive-user-data,performance-tier-inactive-user-data-percent</code> parameter to display how much user data is cold.
Move a volume in to or out of FabricPool	volume move start You use the <code>-tiering-policy</code> optional parameter to specify the tiering policy for the volume.
Modify the threshold for reclaiming unreferenced space (the defragmentation threshold) for FabricPool	storage aggregate object-store modify with the <code>-unreclaimed-space-threshold</code> parameter in the advanced privilege level
Modify the threshold for the percent full the aggregate becomes before the tiering scan begins tiering data for FabricPool	storage aggregate object-store modify with the <code>-tiering-fullness-threshold</code> parameter in the advanced privilege level
Display the threshold for reclaiming unreferenced space for FabricPool	storage aggregate object-store show or storage aggregate object-store show-space command with the <code>-unreclaimed-space-threshold</code> parameter in the advanced privilege level

Managing disks

You can perform various tasks to manage your disks, including removing a failed disk, removing data from a disk, and removing ownership of a disk. There are also tasks you can perform related to managing disks using root-data partitioning and related to self-encrypting disks.

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 and Firmware](#)

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.

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.

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

Displaying 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 using the `storage disk show -ownership` command:

Example

```
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, display the partition ownership using the `storage disk show -partition-ownership` command:

Example

```
cluster::> storage disk show -partition-ownership
Container Container
Disk      Aggregate Root Owner  Owner ID  Data Owner  Owner ID
Owner      Owner ID
-----
1.0.0     -         node1    1886742616 node1    1886742616
node1     1886742616
1.0.1     -         node1    1886742616 node1    1886742616
node1     1886742616
1.0.2     -         node2    1886742657 node2    1886742657
node2     1886742657
1.0.3     -         node2    1886742657 node2    1886742657
node2     1886742657
...
```

Manually assigning ownership of partitioned disks

For partitioned disks, there are three different entities for ownership: the container disk, the data partition, and the root partition. 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. If a container disk fails in a half-populated shelf and is replaced, ONTAP will not auto-assign ownership. In this case, any assignment of new disks will need to be done manually. To make auto-assign work on half-populated shelves, place disks equally on lower half and 6 on far right bays to begin with. That is, 6 disks from bays 0-5 and 6 disks from bays 18-23. After the container disk is assigned in an ADP-

configured system, ONTAP's software will handle any partitioning and partition assignments that are required, without user intervention.

Root-data partitioning

About this task

For root-data partitioning you can think of the three owned entities (the container disk and the two partitions) as being collectively owned by the HA pair. 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 an aggregate, it must be owned by the same node that owns the aggregate.

Steps

1. 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 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>
Data partition	<code>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data true</code>
Root partition	<code>storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -root true</code>

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

Root-data-data partitioning

About this task

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.

Parameters must be used in 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.

Steps

1. 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 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>

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

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 system 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 system 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

Configuring autoassignment of disk ownership

You can configure ONTAP to automatically assign disk ownership according to a disk's stack, shelf, or bay. If configured, automatic disk ownership assignments occur 10 minutes after system initialization and every five minutes during normal system operation.

Before you begin

- Your system must adhere to the requirements for automatic disk ownership.
- If 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.

About this task

The behavior of the **default** autoassignment policy depends on the system model. For entry level models, the **default** policy is equivalent to the **bay** policy. For all other systems, it is equivalent to the **stack** policy.

Steps

1. Configure automatic disk assignment:

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

- 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.

2. Verify the automatic assignment settings for the disks:

```
storage disk option show
```

Example

```
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

Which disk autoassignment policy to use

You can typically use the default autoassignment policy, which is equivalent to the **stack** policy for most systems, and to the **bay** policy for entry-level systems (AFF A2xx, FAS2xxx). However, for some configurations, you might need to change the autoassignment policy.

You must select the appropriate autoassignment based on your configuration:

If you are using...	Then use the autoassignment policy value of...
Stand-alone entry-level system	stack

If you are using...	Then use the autoassignment policy value of...
Entry-level systems in an HA configuration with a single, shared shelf	bay
Entry-level systems in an HA configuration with one stack of two or more shelves	shelf
MetroCluster configurations with one stack per node, two or more shelves	shelf
All other configurations	stack

Removing a failed disk

A disk that is 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. Find the disk ID of the failed disk by entering the following command:

```
storage disk show -broken
```

If the disk does not appear in the list of failed disks, it might be 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 by entering the following command:

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

The fault LED on the face of the disk is lit for 2 minutes.

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

Removing 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.

Before you begin

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 aggregate.
- It cannot be in the maintenance center.

- It cannot be undergoing sanitization.
- It cannot be 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 automatic ownership assignment until the disk is removed, and then reenables it.

Steps

1. If disk ownership automatic assignment is on, turn it off:
`storage disk option modify -node node_name -autoassign off`
2. If needed, repeat the previous step for the node's HA partner.
3. 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, remove ownership from the partitions by entering both of the following commands:
`storage disk removeowner disk_name -root true`
`storage disk removeowner disk_name -data true`
Both partitions are no longer owned by any node.
5. If you turned off disk ownership automatic assignment previously, 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.

Sanitizing a disk

Disk sanitization enables you to remove data from a disk or set of disks so that the data can never be recovered.

Before you begin

- The disk must be a spare, and must be owned by a node but not used in an aggregate.
If the disk is partitioned, neither partition can be used in an aggregate.
- The disk cannot be a self-encrypting disk (SEDs).
You must use the `storage encryption disk sanitize` command to sanitize an SED.
[Encryption of data at rest](#)
- The disk cannot be part of a storage pool.

About this task

When disk sanitization is enabled, it disables some ONTAP commands. After disk sanitization is enabled on a node, it cannot be disabled.

Steps

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

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

2. Enable disk sanitization:

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

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

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

```
disk unpartition disk_name
```

4. Sanitize the specified disks:

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

Attention: 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.

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

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

6. After the sanitization process is complete, return the disks to spare status by entering the following command for each disk:

```
disk sanitize release disk_name
```

7. Exit the nodeshell:

```
exit
```

8. 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.

If...	Then...
Some of the sanitized disks are not listed as spares	<p>Complete the following steps:</p> <ol style="list-style-type: none"> a. Enter advanced privilege mode: <pre>set -privilege advanced</pre> b. Assign the unassigned sanitized disks to the appropriate node by entering the following command for each disk: <pre>storage disk assign -disk <i>disk_name</i> -owner <i>node_name</i></pre> c. Return the disks to spare status by entering the following command for each disk: <pre>storage disk unfail -disk <i>disk_name</i> -s -q</pre> d. Return to administrative mode: <pre>set -privilege admin</pre>

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`.

Setting up an active-passive configuration on nodes using root-data partitioning

When an HA pair is configured to use root-data partitioning by the factory, ownership of the data partitions is split between both nodes in the pair, for use in an active-active configuration. If you want to use the HA pair in an active-passive configuration, you must update partition ownership before creating your data aggregate.

Before you begin

- You should have decided which node will be the active node and which node will be the passive node.
- Storage failover must be configured on the HA pair.

About this task

This task is performed on two nodes: Node A and Node B.

All commands are input at the cluster shell.

This procedure is designed for nodes for which no data aggregate has been created from the partitioned disks.

Steps

1. View the current ownership of the data partitions:

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

Example

You can see 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.

```

cluster1:>> storage aggregate show-spare-disks

Original Owner: cluster1-01
Pool0
  Partitioned Spares

Disk                Type      RPM  Checksum      Local Data Usable  Local Root Usable  Physical 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

Disk                Type      RPM  Checksum      Local Data Usable  Local Root Usable  Physical 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
```

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

```
storage disk assign -force -data true -owner active_node_name -disk disk_name
```

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

Example

You would enter a command similar to the following example for each data partition you need to reassign:

```
storage disk assign -force -data true -owner cluster1-01 -disk 1.0.3
```

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

Example

```

cluster1:>> storage aggregate show-spare-disks

Original Owner: cluster1-01
Pool0
  Partitioned Spares

Disk                Type      RPM  Checksum      Local Data Usable  Local Root Usable  Physical 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

Disk                Type      RPM  Checksum      Local Data Usable  Local Root Usable  Physical Size
-----
1.0.8               BSAS    7200 block        753.8GB  73.89GB 828.0GB
13 entries were displayed.

```

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

- Return to administrative privilege:

```
set admin
```

- 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.

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</code> (advanced)
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 -fields diskpathnames</code>
Display the list of disks in the maintenance center	<code>storage disk show -maintenance center</code>
Display SSD wear life	<code>storage disk show -ssd-wear</code>
Unpartition a disk	<code>system node run -node local -command disk unpartition</code>
Zero all non-zeroed disks	<code>storage disk zerospares</code>
Stop an ongoing sanitization process on one or more specified disks	<code>disk sanitize abort disk_list</code>
Display storage encryption disk information	<code>storage encryption disk show</code>

If you want to...	Use this command...
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(advanced)</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>

If you want to display...	Use this command...
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>
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](#)

Managing RAID groups

You can convert from RAID-DP to RAID-TEC for triple-parity protection, or from RAID-TEC to RAID-DP when a smaller aggregate no longer needs triple-parity protection. You can also customize the size of your RAID groups.

Converting from RAID-DP to RAID-TEC

If you want the added protection of triple-parity, you can convert from RAID-DP to RAID-TEC. RAID-TEC is recommended if the size of the disks used in your aggregate is greater than 4 TiB.

Before you begin

The aggregate that is to be converted must have a minimum of six disks.

About this task

Hard disk drive (HDD) aggregates can be converted from RAID-DP to RAID-TEC. This includes HDD tiers in Flash Pool aggregates.

Steps

1. Verify that the aggregate is online and has a minimum of six disks:

```
storage aggregate show-status -aggregate aggregate_name
```
2. Convert the aggregate from RAID-DP to RAID-TEC:

```
storage aggregate modify -aggregate aggregate_name -raidtype raid_tec
```
3. Verify that the aggregate RAID policy is RAID-TEC:

```
storage aggregate show aggregate_name
```

Converting RAID-TEC to RAID-DP

If you reduce the size of your aggregate and no longer need triple parity, you can convert your RAID policy from RAID-TEC to RAID-DP and reduce the number of disks you need for RAID parity.

About this task

The maximum RAID group size for RAID-TEC is larger than the maximum RAID group size for RAID-DP. If the largest RAID-TEC group size is not within the RAID-DP limits, you cannot convert to RAID-DP.

Steps

1. Verify that the aggregate is online and has a minimum of six disks:

```
storage aggregate show-status -aggregate aggregate_name
```
2. Convert the aggregate from RAID-TEC to RAID-DP:

```
storage aggregate modify -aggregate aggregate_name -raidtype raid_dp
```
3. Verify that the aggregate RAID policy is RAID-DP:

```
storage aggregate show aggregate_name
```

Considerations for sizing RAID groups

Configuring an optimum RAID group size requires a trade-off of factors. You must decide which factors—speed of RAID rebuild, assurance against risk of data loss due to drive failure, optimizing I/O performance, and maximizing data storage space—are most important for the aggregate that you are configuring.

When you create larger RAID groups, you maximize the space available for data storage for the same amount of storage used for parity (also known as the “parity tax”). On the other hand, when a disk fails in a larger RAID group, reconstruction time is increased, impacting performance for a longer period of time. In addition, having more disks in a RAID group increases the probability of a multiple disk failure within the same RAID group.

HDD or array LUN RAID groups

You should follow these guidelines when sizing your RAID groups composed of HDDs or array LUNs:

- All RAID groups in an aggregate should have a similar number of disks.
The RAID groups do not have to be exactly the same size, but you should avoid having any RAID group that is less than one half the size of other RAID groups in the same aggregate when possible.
- The recommended range of RAID group size is between 12 and 20.
The reliability of performance disks can support a RAID group size of up to 28, if needed.
- If you can satisfy the first two guidelines with multiple RAID group sizes, you should choose the larger size.

SSD RAID groups in Flash Pool aggregates

The SSD RAID group size can be different from the RAID group size for the HDD RAID groups in a Flash Pool aggregate. Usually, you should ensure that you have only one SSD RAID group for a Flash Pool aggregate, to minimize the number of SSDs required for parity.

SSD RAID groups in SSD aggregates

You should follow these guidelines when sizing your RAID groups composed of SSDs:

- All RAID groups in an aggregate should have a similar number of drives.
The RAID groups do not have to be exactly the same size, but you should avoid having any RAID group that is less than one half the size of other RAID groups in the same aggregate when possible.
- For RAID-DP, the recommended range of RAID group size is between 20 and 28.

Customizing the size of your RAID groups

You can customize the size of your RAID groups to ensure that your RAID group sizes are appropriate for the amount of storage you plan to include for an aggregate.

About this task

For standard aggregates, you change the size of RAID groups on a per-aggregate basis. For Flash Pool aggregates, you can change the RAID group size for the SSD RAID groups and the HDD RAID groups independently.

The following list outlines some facts about changing the RAID group size:

- By default, if the number of disks or array LUNs in the most recently created RAID group is less than the new RAID group size, disks or array LUNs will be added to the most recently created RAID group until it reaches the new size.
- All other existing RAID groups in that aggregate remain the same size, unless you explicitly add disks to them.
- You can never cause a RAID group to become larger than the current maximum RAID group size for the aggregate.
- You cannot decrease the size of already created RAID groups.
- The new size applies to all RAID groups in that aggregate (or, in the case of a Flash Pool aggregate, all RAID groups for the affected RAID group type—SSD or HDD).

Step

1. Use the applicable command:

If you want to...	Enter the following command...
Change the maximum RAID group size for the SSD RAID groups of a Flash Pool aggregate	<code>storage aggregate modify -aggregate <i>aggr_name</i> -cache-raid-group-size <i>size</i></code>
Change the maximum size of any other RAID groups	<code>storage aggregate modify -aggregate <i>aggr_name</i> -maxraidsize <i>size</i></code>

Examples

The following command changes the maximum RAID group size of the aggregate `n1_a4` to 20 disks or array LUNs:

```
storage aggregate modify -aggregate n1_a4 -maxraidsize 20
```

The following command changes the maximum RAID group size of the SSD cache RAID groups of the Flash Pool aggregate `n1_cache_a2` to 24:

```
storage aggregate modify -aggregate n1_cache_a2 -cache-raid-group-size 24
```

Related concepts

[Considerations for sizing RAID groups](#) on page 74

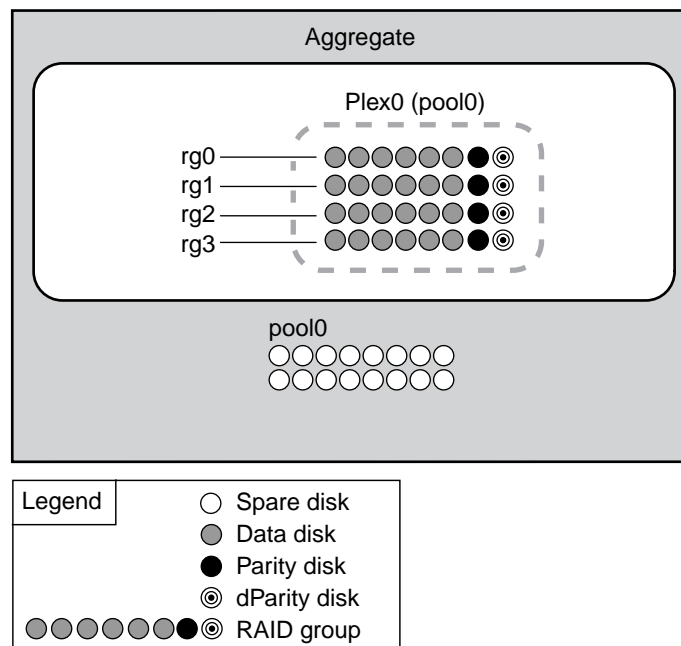
Appendix: Mirrored and unmirrored aggregates

You can use an optional feature called *SyncMirror* to synchronously mirror aggregate data in copies, or *plexes*, stored in different RAID groups. Plexes ensure against data loss if more disks fail than the RAID type protects against, or if there is a loss of connectivity to RAID group disks.

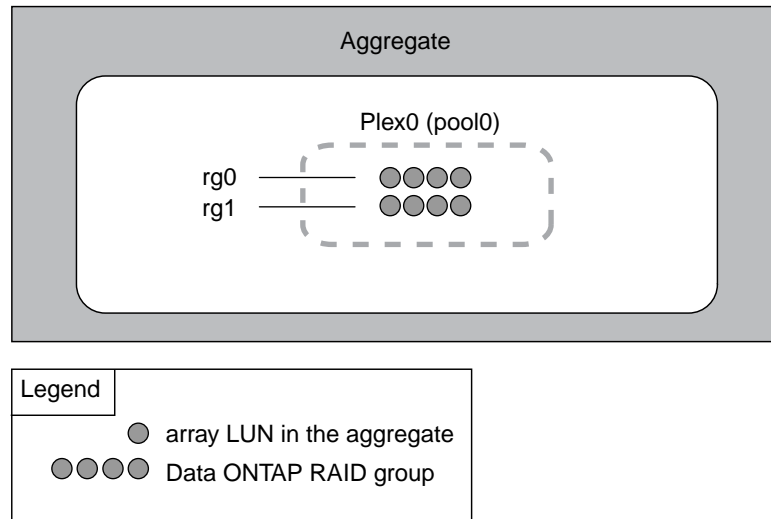
How unmirrored aggregates work

Unless you are using *SyncMirror*, all of your aggregates are unmirrored. Unmirrored aggregates have only one *plex* (copy of their data), which contains all of the RAID groups belonging to that aggregate.

The following diagram shows an unmirrored aggregate composed of disks, with its one plex. The aggregate has four RAID groups: rg0, rg1, rg2, and rg3. Each RAID group has 6 data disks, one parity disk, and one dparity (double parity) disk. All disks used by the aggregate come from the same pool, pool0.



The following diagram shows an unmirrored aggregate with array LUNs, with its one plex. It has two RAID groups, rg0 and rg1. All array LUNs used by the aggregate come from the same pool, pool0.



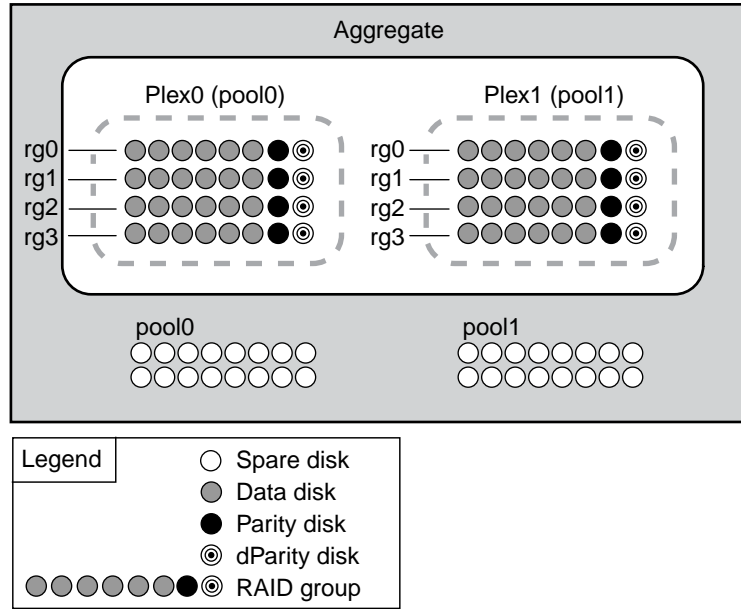
How mirrored aggregates work

Mirrored aggregates have two *plexes* (copies of their data), which use the SyncMirror functionality to duplicate the data to provide redundancy.


When a mirrored aggregate is created (or when a second plex is added to an existing unmirrored aggregate), ONTAP copies the data in the original plex (plex0) to the new plex (plex1). The plexes are physically separated (each plex has its own RAID groups and its own pool), and the plexes are updated simultaneously. This provides added protection against data loss if more disks fail than the RAID level of the aggregate protects against or there is a loss of connectivity, because the unaffected plex continues to serve data while you fix the cause of the failure. After the plex that had a problem is fixed, the two plexes resynchronize and reestablish the mirror relationship.

The disks and array LUNs on the system are divided into two pools: pool0 and pool1. Plex0 gets its storage from pool0 and plex1 gets its storage from pool1.

The following diagram shows an aggregate composed of disks with the SyncMirror functionality enabled and implemented. A second plex has been created for the aggregate, plex1. The data in plex1 is a copy of the data in plex0, and the RAID groups are also identical. The 32 spare disks are allocated to pool0 or pool1, 16 disks for each pool.



The following diagram shows an aggregate composed of array LUNs with the SyncMirror functionality enabled and implemented. A second plex has been created for the aggregate, plex1. Plex1 is a copy of plex0, and the RAID groups are also identical.



Where to find additional information

After you have setup the disks and aggregates for your storage system, you can perform additional tasks such as setting up volumes to contain your data.

- [*ONTAP 9 commands*](#)
Describe commands for creating and managing aggregates in reference format.
- [*Logical storage management*](#)
Describes how to manage logical storage resources in clusters, including FlexVol volumes, FlexClone volumes, FlexCache volumes, files, and LUNs.
- [*SAN administration*](#)
Describes how to configure and manage iSCSI, FCoE, and FC protocols including configuration of LUNs, igroups and targets.
- [*NetApp Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment*](#)
Describes thin provisioning in a SAN or IP SAN enterprise environment.
- [*NFS management*](#)
Describes how to manage file access to clusters using the NFS protocol, including authentication, authorization, and security
- [*SMB/CIFS management*](#)
Describes how to configure CIFS servers and manage file access using the SMB protocol, including security, file policy management, and CIFS client- and server-based services.

Copyright

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- about this guide
 - deciding whether to use the Disks and Aggregates Power Guide [6](#)
 - active-passive configurations
 - setting up with root-data partitioning [68](#)
 - additional information
 - where to find [79](#)
 - aggregate creation
 - deciding which method to use [7](#)
 - workflow [7](#)
 - aggregate expansion
 - workflow [12](#)
 - aggregate relocation
 - commands for [22](#)
 - aggregate show-space command
 - how to determine aggregate space usage by using [23](#)
 - aggregates
 - assigning to SVMs [22](#)
 - attaching to the cloud tier of FabricPool [50](#)
 - changing size of RAID groups for [74](#)
 - choices for adding or moving volumes to FabricPool [53](#)
 - commands for displaying space usage information [71](#)
 - commands for managing [26](#)
 - commands for managing with FabricPool [58](#)
 - considerations for sizing RAID groups [74](#)
 - creating Flash Pool [31](#)
 - creating Flash Pool using SSD storage pools [33](#)
 - creating manually [10](#)
 - creating with auto-provision [8](#)
 - default RAID policies for [9](#)
 - determining candidacy and optimal cache size for Flash Pool [29](#)
 - determining drive information for [19](#)
 - determining RAID group information for [19](#)
 - determining which volumes reside on [24](#)
 - expanding [14](#)
 - expansion workflow [12](#)
 - Flash Pool, determining whether using a storage pool [32](#)
 - how to determine and control volume space usage in [24](#)
 - how to determine space usage in [23](#)
 - how to determine the number of disks required for [9](#), [17](#)
 - introduction to managing [17](#)
 - introduction to relocating ownership within an HA pair [20](#)
 - managing storage tiers in ONTAP 9.2 or later by using FabricPool, introduction [39](#)
 - methods of creating space in [25](#)
 - mirrored, explained [77](#)
 - moving a volume to FabricPool [54](#)
 - ownership change [20](#)
 - relocation of [20](#)
 - requirements for using the Disks and Aggregates Power Guide [6](#)
 - unmirrored, explained [76](#)
 - where to find additional information about [79](#)
 - aggregates, Flash Pool
 - creating SSD storage pools for [32](#)
 - how caching policies work [27](#)
 - modifying caching policies of [36](#)
 - setting cache-retention policy for [37](#)
 - aggregates, Flash Pool using storage pools
 - how Flash Pool SSD partitioning works for [27](#)
 - all flash aggregates
 - how to determine the number of disks required for [9](#), [17](#)
 - array LUNs
 - considerations for sizing RAID groups for [74](#)
 - assigning ownership
 - manually for partitioned disks [61](#)
 - manually for unpartitioned disks [13](#)
 - auto-provision
 - creating aggregates with [8](#)
 - autoassignment, of disk ownership
 - configuring [64](#)
 - Automated Workload Analyzer
 - determining Flash Pool optimal cache size by using [29](#)
 - automatic ownership assignment
 - which policy to use for disks [64](#)
 - AWA
 - determining Flash Pool optimal cache size by using [29](#)
 - AWS S3
 - installing a FabricPool license if required for [45](#)
 - setting up as the cloud tier [47](#)
 - Azure Blob Storage
 - installing a FabricPool license if required for [45](#)
 - setting up as the cloud tier [48](#)
- ## C
- CA certificates
 - installing if you use StorageGRID as the object store for FabricPool [46](#)
 - cache size
 - determining impact to, when adding SSDs to storage pool [34](#)
 - determining optimal for Flash Pool aggregates [29](#)
 - cache-retention policies
 - setting for Flash Pool aggregates [37](#)
 - caching policies
 - determining whether to modify for Flash Pool aggregates [36](#)
 - how they work for Flash Pool aggregates [27](#)
 - modifying, of Flash Pool aggregates [36](#)
 - capacity HDD aggregates
 - how to determine the number of disks required for [9](#), [17](#)
 - changing
 - RAID group size [74](#)

- cloud tier
 - external, attaching AWS S3 as a FabricPool using OnCommand System Manager (video) [48](#)
- cloud tiers
 - attaching to an aggregate for FabricPool [50](#)
 - external, attaching StorageGRID as a FabricPool using OnCommand System Manager (video) [50](#)
 - setting up AWS S3 [47](#)
 - setting up Azure Blob Storage [48](#)
 - setting up IBM Cloud Object Storage [49](#)
 - setting up StorageGRID [49](#)
 - types of object stores that you can use for FabricPool [47](#)
- commands
 - aggregate management, list of [26](#)
 - disk management, list of [70](#)
 - for displaying aggregate space usage information [71](#)
 - for displaying FlexVol volume space usage information [71](#)
 - for displaying information about storage shelves [71](#)
 - for managing aggregates with FabricPool [58](#)
 - SSD storage pool management, list of [35](#)
 - storage aggregate [26](#), [70](#), [71](#)
 - storage disk [70](#)
 - storage pool [35](#)
 - storage pool management, list of [35](#)
 - storage shelves, displaying information about [71](#)
 - volume show-footprint [71](#)
 - volume show-space [71](#)
 - volume snapshot [71](#)
- comments
 - how to send feedback about documentation [82](#)
- configuring
 - autoassignment of disk ownership [64](#)
- considerations
 - and requirements for using FabricPool [39](#)
- constituents
 - determining which ones reside on an aggregate [24](#)
- converting
 - RAID-DP to RAID-TEC [73](#)
 - RAID-TEC to RAID-DP [73](#)
- creating
 - aggregates manually [10](#)
 - Flash Pool aggregates [31](#)
 - Flash Pool aggregates using SSD storage pools [33](#)
- creating aggregates
 - workflow [7](#)

D

- data
 - using sanitization to remove from disks [66](#)
- deciding which method to use
 - aggregate creation [7](#)
- disk ownership
 - configuring autoassignment of [64](#)
- Disk Qualification Package
 - when you need to update [60](#)
- disk shelves
 - commands for displaying information about [71](#)
- disks
 - adding to a node [13](#)
 - commands for managing [70](#)

- determining information about, for aggregates [19](#)
- how low spare warnings can help you manage spare [60](#)
- how to determine the number required for an aggregate [9](#), [17](#)
- overview of managing [60](#)
- RAID protection levels for [17](#)
- removing data from using disk sanitization [66](#)
- removing failed [65](#)
- removing ownership from [65](#)
- requirements for using the Disks and Aggregates Power Guide [6](#)
- viewing ownership for [61](#)
- when you need to update the Disk Qualification Package for [60](#)
- which autoassignment policy to use for [64](#)

Disks and Aggregates Power Guide requirements for using [6](#)

disks, hot spares

- how they work [60](#)

disks, partitioned

- manually assigning ownership for [61](#)

disks, unpartitioned

- manually assigning ownership for [13](#)

documentation

- how to receive automatic notification of changes to [82](#)
- how to send feedback about [82](#)
- where to find additional information [79](#)

DQP

- Disk Qualification Package [60](#)

drives

- considerations for sizing RAID groups for [74](#)
- how low spare warnings can help you manage spare [60](#)

E

- expanding
 - aggregates [14](#)
- expanding aggregates
 - workflow [12](#)

F

FabricPools

- archiving volumes with using OnCommand System Manager (video) [57](#)
- attaching AWS S3 as a cloud tier using OnCommand System Manager (video) [48](#)
- attaching StorageGRID as a cloud tier using OnCommand System Manager (video) [50](#)
- attaching the cloud tier to an aggregate [50](#)
- benefits of storage tiers by using [39](#)
- choices for adding or moving volumes to [53](#)
- commands for managing aggregates with [58](#)
- configuration and management workflow diagram [44](#)
- configuring [44](#)
- considerations and requirements for using [39](#)
- creating a volume for [53](#)
- determining how much data in a volume is inactive by using inactive data reporting [51](#)

- installing a CA certificate if you use StorageGRID as the object store [46](#)
 - installing a license if required for AWS S3, Azure Blob Storage, or IBM Cloud Object Storage [45](#)
 - installing a license using OnCommand System Manager (video) [46](#)
 - introduction to managing [51](#)
 - managing storage tiers in ONTAP 9.2 or later by using, introduction to [39](#)
 - modifying the tiering policy of a volume for [56](#)
 - monitoring the space utilization for [55](#)
 - moving a volume to [54](#)
 - preparing for configuration [45](#)
 - setting up AWS S3 as the cloud tier [47](#)
 - setting up Azure Blob Storage as the cloud tier [48](#)
 - setting up IBM Cloud Object Storage as the cloud tier [49](#)
 - setting up StorageGRID as the cloud tier [49](#)
 - types of object stores that you can use as the cloud tier [47](#)
 - understanding tiering policies [42](#)
 - failed disks
 - removing [65](#)
 - feedback
 - how to send comments about documentation [82](#)
 - Flash Pool aggregates
 - creating [31](#)
 - creating SSD storage pools for [32](#)
 - creating using SSD storage pools [33](#)
 - determining candidacy and optimal cache size for [29](#)
 - determining whether to modify caching policy of [36](#)
 - determining whether using a storage pool [32](#)
 - how caching policies work [27](#)
 - how to determine the number of disks required for [9](#), [17](#)
 - modifying caching policies of [36](#)
 - setting the cache-retention policy for [37](#)
 - Flash Pool aggregates using storage pools
 - how Flash Pool SSD partitioning works for [27](#)
 - Flash Pool SSD partitioning
 - how it works for Flash Pool aggregates using storage pools [27](#)
 - FlexVol volumes
 - commands for displaying space usage information [71](#)
 - determining which ones reside on an aggregate [24](#)
 - how to determine and control space usage in aggregates [24](#)
- ## H
- HA pairs
 - introduction to relocating aggregate ownership within [20](#)
 - HDD RAID groups
 - sizing considerations for [74](#)
 - hot spare disks
 - how they work [60](#)
- ## I
- IBM Cloud Object Storage
 - installing a FabricPool license if required for [45](#)
 - setting up as the cloud tier [49](#)
 - inactive data reporting
 - determining how much data in a volume is inactive by using [51](#)
 - Infinite Volumes
 - determining which constituents reside on an aggregate [24](#)
 - how to determine and control space usage for constituents in aggregates [24](#)
 - information
 - how to send feedback about improving documentation [82](#)
- ## L
- licenses
 - installing FabricPool using OnCommand System Manager (video) [46](#)
 - installing for FabricPool, if required for AWS S3, Azure Blob Storage, or IBM Cloud Object Storage [45](#)
 - low spare warnings
 - how they can help you manage spare drives [60](#)
- ## M
- managing
 - disks, overview of [60](#)
 - root-data partitions [63](#)
 - mirrored aggregates
 - explained [77](#)
 - modifying
 - caching policies of Flash Pool aggregates [36](#)
- ## N
- nodes
 - adding disks to [13](#)
 - introduction to relocating aggregate ownership within an HA pair [20](#)
- ## O
- object stores
 - types that you can use as the cloud tier for FabricPool [47](#)
 - overviews
 - of managing disks [60](#)
 - ownership
 - displaying disk ownership [61](#)
 - displaying partition [61](#)
 - manually assigning for unpartitioned disks [13](#)
 - removing from disks [65](#)
 - ownership, aggregate
 - introduction to relocating within an HA pair [20](#)
 - ownership, disk
 - configuring autoassignment of [64](#)
 - manually assigning for partitioned disks [61](#)

P

- parity disks
 - how to determine the number required for an aggregate [9](#), [17](#)
- partitioned disks
 - manually assigning ownership for [61](#)
- partitioning
 - setting up active-passive configuration on node using root-data [68](#)
- partitioning, Flash Pool SSD
 - how it works for Flash Pool aggregates using storage pools [27](#)
- partitions
 - correcting misaligned spare [18](#)
 - viewing ownership for [61](#)
- performance HDD aggregates
 - how to determine the number of disks required for [9](#), [17](#)
- plexes
 - mirrored aggregate, explained [77](#)
- policies
 - modifying caching of Flash Pool aggregates [36](#)
 - which to use for disk autoassignment [64](#)
- policies, cache-retention
 - setting for Flash Pool aggregates [37](#)
- policies, caching
 - how they work for Flash Pool aggregates [27](#)
- pools, SSD storage
 - adding SSDs to [35](#)
 - creating [32](#)
- power guides
 - aggregate creation workflow [7](#)
 - aggregate expansion workflow [12](#)
 - requirements for using the Disks and Aggregates Power Guide [6](#)
- protection levels, RAID
 - for disks [17](#)

R

- RAID
 - protection levels for disks [17](#)
- RAID groups
 - changing size of [74](#)
 - determining information about, for aggregates [19](#)
 - sizing considerations for [74](#)
- RAID policies
 - default for aggregates [9](#)
- RAID-DP
 - converting to RAID-TEC [73](#)
 - converting to, from RAID-TEC [73](#)
- RAID-TEC
 - converting to RAID-DP [73](#)
 - converting to, from RAID-DP [73](#)
- relocating
 - aggregate ownership within an HA pair, introduction to [20](#)
- relocation
 - aggregate ownership [20](#)
 - of aggregates [20](#)
- removing
 - data from disk, using disk sanitization [66](#)

- failed disks [65](#)
- requirements
 - and considerations for using FabricPool [39](#)
- root-data partitioning
 - additional management options [63](#)
 - correcting misaligned spare partitions for [18](#)
 - setting up active-passive configuration on node using [68](#)

S

- sanitization
 - using to remove data from disks [66](#)
- setting
 - cache-retention policy for Flash Pool aggregates [37](#)
- shelves
 - commands for displaying information about [71](#)
- sizing
 - RAID groups, considerations for [74](#)
- Snapshot reserve
 - commands for displaying size of [71](#)
- space
 - commands for displaying usage information [71](#)
 - methods of creating in an aggregate [25](#)
- space usage
 - how to determine and control volume, in aggregates [24](#)
 - how to determine in an aggregate [23](#)
- space utilization
 - monitoring for FabricPool [55](#)
- spare disks
 - how low spare warnings can help you manage [60](#)
 - how to determine the number required for an aggregate [9](#), [17](#)
 - removing ownership from [65](#)
- spare partitions
 - correcting misaligned [18](#)
- SSD storage pools
 - commands for managing [35](#)
 - creating [32](#)
 - creating Flash Pool aggregates using [33](#)
 - how Flash Pool SSD partitioning works for Flash Pool aggregates using [27](#)
- SSDs
 - adding to SSD storage pools [35](#)
 - changing size of RAID groups for [74](#)
 - determining impact to cache size of adding to storage pool [34](#)
 - sizing considerations for RAID groups [74](#)
 - storage pools, determining when used by a Flash Pool aggregate [32](#)
- storage aggregate commands
 - for displaying space information [71](#)
 - for managing aggregates [26](#)
 - for managing disks [70](#)
- storage disk commands
 - for managing disks [70](#)
- storage pools
 - commands for managing [35](#)
 - creating Flash Pool aggregates using SSD [33](#)
 - determining impact to cache size of adding SSDs to [34](#)
 - determining when used by a Flash Pool aggregate [32](#)

- storage pools, SSD
 - adding SSDs to [35](#)
 - creating [32](#)
 - how Flash Pool SSD partitioning works for Flash Pool aggregates using [27](#)
- storage shelves
 - commands for displaying information about [71](#)
- storage tiers
 - benefits of, by using FabricPool [39](#)
 - managing in ONTAP 9.2 or later by using FabricPool, introduction to [39](#)
- StorageGRID
 - installing a CA certificate if you use as the object store for FabricPool [46](#)
 - setting up as the cloud tier [49](#)
- suggestions
 - how to send feedback about documentation [82](#)
- SVMs
 - assigning aggregates to [22](#)

T

- tiering policies
 - creating a volume for FabricPool [53](#)
 - modifying for a volume for FabricPool [56](#)
 - understanding FabricPool [42](#)
- Twitter
 - how to receive automatic notification of documentation changes [82](#)

U

- unmirrored aggregates
 - explained [76](#)
- unpartitioned disks
 - manually assigning ownership for [13](#)
- used space
 - how to determine and control in aggregates, by volume [24](#)

- how to determine in aggregate [23](#)

V

- videos
 - archiving volumes with FabricPool using OnCommand System Manager [57](#)
 - attaching AWS S3 as a FabricPool cloud tier using OnCommand System Manager [48](#)
 - attaching StorageGRID as a FabricPool cloud tier using OnCommand System Manager [50](#)
 - installing a FabricPool license [46](#)
- volume command
 - for displaying space information [71](#)
- volume show-footprint command
 - understanding output [24](#)
- volumes
 - archiving with FabricPool using OnCommand System Manager (video) [57](#)
 - choices for adding or moving to FabricPool [53](#)
 - creating for FabricPool [53](#)
 - determining which ones reside on an aggregate [24](#)
 - how caching policies work for Flash Pool aggregates [27](#)
 - how to determine space usage of, in aggregates [24](#)
 - modifying the tiering policy for FabricPool [56](#)
 - moving to FabricPool [54](#)
 - where to find additional information about [79](#)
- Vservers
 - See* SVMs

W

- workflows
 - aggregate creation [7](#)
 - aggregate expansion [12](#)