



Storage tiers

System Manager Classic

NetApp
January 13, 2022

Table of Contents

- Storage tiers 1
 - Editing aggregates 1
 - Deleting aggregates 2
 - Changing the RAID configuration when creating an aggregate 2
 - Provisioning cache by adding SSDs 3
 - Add capacity disks 5
 - Changing the RAID group when adding capacity disks 7
 - Moving FlexVol volumes 7
 - Mirroring aggregates 8
 - Viewing aggregate information 8
 - Install a CA certificate if you use StorageGRID 9
 - How moving a FlexVol volume works 9
 - How you can use effective ONTAP disk type for mixing HDDs 10
 - What compatible spare disks are 10
 - How System Manager works with hot spares 11
 - Rules for displaying disk types and disk RPM 12
 - How mirrored aggregates work 12
 - What a FabricPool is 13
 - Storage recommendations for creating aggregates 14
 - Storage Tiers window 14

Storage tiers

You can use System Manager to create aggregates to support the different security requirements, backup requirements, performance requirements, and data sharing requirements of your users.

Related information

[Disk and aggregate management](#)

Editing aggregates

You can use System Manager to change the aggregate name, RAID type, and RAID group size of an existing aggregate when required.

Before you begin

For modifying the RAID type of an aggregate from RAID4 to RAID-DP, the aggregate must contain enough compatible spare disks, excluding the hot spares.

About this task

- You cannot change the RAID group of ONTAP systems that support array LUNs.
RAID0 is the only available option.
- You cannot change the RAID type of partitioned disks.
RAID-DP is the only option that is available for partitioned disks.
- You cannot rename a SnapLock Compliance aggregate.
- If the aggregate consists of SSDs with storage pool, you can modify only the name of the aggregate.
- If the triple parity disk size is 10 TB, and the other disks are smaller than 10 TB in size, then you can select RAID-DP or RAID-TEC as the RAID type.
- If the triple parity disk size is 10 TB, and if even one of the other disks is larger than 10 TB in size, then RAID-TEC is the only available option for RAID type.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. Select the aggregate that you want to edit, and then click **Edit**.
3. In the **Edit Aggregate** dialog box, modify the aggregate name, the RAID type, and the RAID group size, as required.
4. Click **Save**.

Related information

[Aggregates window](#)

Deleting aggregates

You can use System Manager to delete aggregates when you no longer require the data in the aggregates. However, you cannot delete the root aggregate because it contains the root volume, which contains the system configuration information.

Before you begin

- All the FlexVol volumes and the associated storage virtual machines (SVMs) contained by the aggregate must be deleted.
- The aggregate must be offline.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. Select one or more aggregates that you want to delete, and then click **Delete**.
3. Select the confirmation check box, and then click **Delete**.

Related information

Changing the RAID configuration when creating an aggregate

While creating an aggregate, you can modify the default values of the RAID type and RAID group size options of the aggregate by using System Manager.

About this task

If the disk type of the aggregate disks is FSAS or MSATA, and the disk size is equal to or larger than 10 TB, then RAID-TEC is the only available RAID type.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. In the **Storage Tiers** window, click **Add Aggregate**.
3. In the **Create Aggregate** dialog box, perform the following steps:
 - a. Click **Change**.
 - b. In the **Change RAID Configuration** dialog box, specify the RAID type and RAID group size.

Shared disks support two RAID types: RAID DP and RAID-TEC.

The recommended RAID group size is 12 disks through 20 disks for HDDs, and 20 disks through 28 disks for SSDs.

c. Click **Save**.

Provisioning cache by adding SSDs

You can use System Manager to add SSDs as either storage pools or dedicated SSDs to provision cache. By adding SSDs, you can convert a non-root aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate, or increase the cache size of an existing Flash Pool aggregate.

About this task

- The added SSD cache does not add to the size of the aggregate, and you can add an SSD RAID group to an aggregate even when it is at the maximum size.
- You cannot use partitioned SSDs when you add cache by using System Manager.

Related information

[How storage pool works](#)

Provisioning cache to aggregates by adding SSDs

You can use System Manager to add storage pools or dedicated SSDs to provision cache by converting an existing non-root HDD aggregate or a root aggregate that does not contain partitioned disks to a Flash Pool aggregate.

Before you begin

- The aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.
- All of the nodes in the cluster must be running ONTAP 8.3 or later.

If the cluster is in a mixed-version state, you can use the command-line interface to create a Flash Pool aggregate and then provision SSD cache.

- You must have identified a valid 64-bit non-root aggregate composed of HDDs that can be converted to a Flash Pool aggregate.
- The aggregate must not contain any array LUNs.

About this task

You must be aware of platform-specific and workload-specific best practices for Flash Pool aggregate SSD tier size and configuration.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.

- Click **Storage > Aggregates & Disks > Aggregates**.

2. In the **Storage Tiers** window, select the aggregate, and then click **More Actions > Add Cache**.



Adding cache is not supported on FabricPool-enabled aggregates.

3. In the **Add Cache** dialog box, perform the appropriate action:

If you selected the cache source as...	Do this...
Storage pools	a. Select the storage pool from which cache can be obtained. b. Specify the cache size. c. Modify the RAID type, if required.
Dedicated SSDs	Select the SSD size and the number of SSDs to include, and optionally modify the RAID configuration: a. Click Change . b. In the Change RAID Configuration dialog box, specify the RAID type and RAID group size, and then click Save .

4. Click **Add**.

For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.

5. In the **Add Cache** dialog box, click **Yes**.

Results

The cache disks are added to the selected aggregate.

Related information

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

Increasing the cache for Flash Pool aggregates by adding SSDs

You can add SSDs as either storage pools or dedicated SSDs to increase the size of a Flash Pool aggregate by using System Manager.

Before you begin

- The Flash Pool aggregate must be online.
- There must be sufficient spare SSDs or allocation units in the storage pool that can be assigned as cache disks.

Steps

1. Click **Storage > Aggregates & Disks > Aggregates**.

2. In the **Aggregates** window, select the Flash Pool aggregate, and then click **Add Cache**.
3. In the **Add Cache** dialog box, perform the appropriate action:

If you selected the cache source as...	Do this...
Storage pools	Select the storage pool from which cache can be obtained, and specify the cache size.
Dedicated SSDs	Select the SSD size and the number of SSDs to include.

4. Click **Add**.

For mirrored aggregates, an Add Cache dialog box is displayed with the information that twice the number of selected disks will be added.

5. In the **Add Cache** dialog box, click **Yes**.

Results

The cache disks are added to the selected Flash Pool aggregate.

Add capacity disks

You can increase the size of an existing non-root aggregate or a root aggregate containing disks by adding capacity disks. You can use System Manager to add HDDs or SSDs of the selected ONTAP disk type and to modify the RAID group options.

Before you begin

- The aggregate must be online.
- There must be sufficient compatible spare disks.

About this task

- It is a best practice to add disks that are of the same size as the other disks in the aggregate.

If you add disks that are smaller in size than the other disks in the aggregate, the aggregate becomes suboptimal in configuration, which in turn might cause performance issues.

If you add disks that are larger in size than the disks that are available in a pre-existing RAID group within the aggregate, then the disks are downsized, and their space is reduced to that of the other disks in that RAID group. If a new RAID group is created in the aggregate and similar sized disks remain in the new RAID group, the disks are not downsized.

If you add disks that are not of the same size as the other disks in the aggregate, the selected disks might not be added; instead, other disks with a usable size between 90 percent and 105 percent of the specified size are automatically added. For example, for a 744 GB disk, all of the disks in the range of 669 GB through 781 GB are eligible for selection. For all of the spare disks in this range, ONTAP first selects only partitioned disks, then selects only unpartitioned disks, and finally selects both partitioned disks and unpartitioned disks.

- You cannot use System Manager to add HDDs to the following configurations:

- Aggregates containing only SSDs
- Root aggregates containing partitioned disks You must use the command-line interface to add HDDs to these configurations.
- Shared disks support two RAID types: RAID DP and RAID-TEC.
- You cannot use SSDs with storage pool.
- If the RAID group type is RAID DP, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them only to `Specific RAID group`, and not to `New RAID group` or `All RAID groups`.

The disks are added after downsizing the disk size to the size of the disks in the pre-existing RAID group of the existing aggregate.

- If the RAID group type is RAID-TEC, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them to `All RAID groups`, `New RAID group`, and `Specific RAID group`.

The disks are added after downsizing the disk size to the size of the disks in the pre-existing RAID group of the existing aggregate.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. In the **Storage Tiers** window, select the aggregate to which you want to add capacity disks, and then click **More Actions > Add Capacity**.
3. Specify the following information in the **Add Capacity** dialog box:
 - a. Specify the disk type for the capacity disks by using the **Disk Type to Add** option.
 - b. Specify the number of capacity disks by using the **Number of Disks or Partitions** option.
4. Specify the RAID group to which the capacity disks are to be added by using the **Add Disks To** option.

By default, System Manager adds the capacity disks to `All RAID groups`.

- a. Click **Change**.
- b. In the **RAID Group Selection** dialog box, specify the RAID group as `New RAID group` or `Specific RAID group` by using the **Add Disks To** option.

Shared disks can be added only to the `New RAID group` option.

5. Click **Add**.

For mirrored aggregates, an Add Capacity dialog box is displayed with the information that twice the number of selected disks will be added.

6. In the **Add Capacity** dialog box, click **Yes** to add the capacity disks.

Results

The capacity disks are added to the selected aggregate, and the aggregate size is increased.

Related information

[What compatible spare disks are](#)

Changing the RAID group when adding capacity disks

While adding capacity disks (HDDs) to an aggregate, you can change the RAID group to which you want to add the disks by using System Manager.

About this task

- If the RAID type is RAID-DP, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them only to `Specific RAID group`, and not to `New RAID group` or `All RAID groups`.

The disks are added after downsizing the disk size to the size of the existing aggregates.

- If the RAID group is RAID-TEC, and if you are adding FSAS or MSATA type of disks that are equal to or larger than 10 TB in size, then you can add them to `All RAID groups`, `New RAID group`, and `Specific RAID group`.

The disks are added after downsizing the disk size to the size of the existing aggregates.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. In the **Storage Tiers** window, select the aggregate to which you want to add capacity disks, and then click **More Actions > Add Capacity**.
3. In the **Add Capacity** dialog box, perform the following steps:
 - a. Click **Change**.
 - b. In the **Change RAID Configuration** dialog box, specify the RAID group to which you want to add the capacity disks.

You can change the default value `All RAID groups` to either `Specific RAID group` or `New RAID group`.
 - c. Click **Save**.

Moving FlexVol volumes

You can nondisruptively move a FlexVol volume to a different aggregate or a different node for capacity utilization and improved performance by using System Manager.

Before you begin

If you are moving a data protection volume, data protection mirror relationships must be initialized before you move the volume.

About this task

- When you move a volume that is hosted on a Flash Pool aggregate, only the data that is stored in the HDD tier is moved to the destination aggregate.

The cached data that is associated with the volume is not moved to the destination aggregate. Therefore, some performance degradation might occur after the volume move.

- You cannot move volumes from a SnapLock aggregate.
- You cannot move volumes from an SVM that is configured for disaster recovery to a FabricPool-enabled aggregate.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. Select the aggregate that contains the volume, and then click **More Actions > Volume Move**.
3. Type or select information as prompted by the wizard.
4. Confirm the details, and then click **Finish** to complete the wizard.

Mirroring aggregates

You can use System Manager to protect data and to provide increased resiliency by mirroring data in real-time, within a single aggregate. Mirroring aggregates removes single points of failure in connecting to disks and array LUNs.

Before you begin

There must be sufficient free disks in the other pool to mirror the aggregate.

About this task

You cannot mirror a Flash Pool aggregate when the cache source is storage pool.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. Select the aggregate that you want to mirror, and then click **More Actions > Mirror**.



SyncMirror is not supported on FabricPool-enabled aggregates.

3. In the **Mirror this aggregate** dialog box, click **Mirror** to initiate the mirroring.

Viewing aggregate information

You can use the Aggregates window in System Manager to view the name, status, and space information about an aggregate.

Steps

1. Choose one of the following methods:
 - Click **Applications & Tiers > Storage Tiers**.
 - Click **Storage > Aggregates & Disks > Aggregates**.
2. Click on the aggregate name to view the details of the selected aggregate.

Install a CA certificate if you use StorageGRID

For ONTAP to authenticate with StorageGRID as the object store for a FabricPool-enabled aggregate, you can install a StorageGRID CA certificate on the cluster.

Steps

1. Follow the StorageGRID system documentation to copy the CA certificate of the StorageGRID system by using the Grid Management Interface.

[StorageGRID 11.3 Administrator Guide](#)

While adding StorageGRID as a cloud tier, a message is displayed if the CA certificate is not installed.

2. Add the StorageGRID CA certificate.



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

Related information

[Adding a cloud tier](#)

How moving a FlexVol volume works

Knowing how moving a FlexVol volume works helps you to determine whether the volume move satisfies service-level agreements and to understand where a volume move is in the volume move process.

FlexVol volumes are moved from one aggregate or node to another within the same storage virtual machine (SVM). A volume move does not disrupt client access during the move.

Moving a volume occurs in multiple phases:

- A new volume is made on the destination aggregate.
- The data from the original volume is copied to the new volume.

During this time, the original volume is intact and available for clients to access.

- At the end of the move process, client access is temporarily blocked.

During this time the system performs a final replication from the source volume to the destination volume, swaps the identities of the source and destination volumes, and changes the destination volume to the source volume.

- After completing the move, the system routes client traffic to the new source volume and resumes client access.

The move is not disruptive to client access because the time in which client access is blocked ends before clients notice a disruption and time out. Client access is blocked for 35 seconds by default. If the volume move operation cannot finish in the time that access is denied, the system aborts this final phase of the volume move operation and allows client access. The system attempts the final phase three times by default. After the third attempt, the system waits an hour before attempting the final phase sequence again. The system runs the final phase of the volume move operation until the volume move is complete.

How you can use effective ONTAP disk type for mixing HDDs

Starting with Data ONTAP 8.1, certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and managing spares. ONTAP assigns an effective disk type for each disk type. You can mix HDDs that have the same effective disk type.

When the `raid.disktype.enable` option is set to `off`, you can mix certain types of HDDs within the same aggregate. When the `raid.disktype.enable` option is set to `on`, the effective disk type is the same as the ONTAP disk type. Aggregates can be created using only one disk type. The default value for the `raid.disktype.enable` option is `off`.

Starting with Data ONTAP 8.2, the option `raid.mix.hdd.disktype.capacity` must be set to `on` to mix disks of type BSAS, FSAS, and ATA. The option `raid.mix.hdd.disktype.performance` must be set to `on` to mix disks of type FCAL and SAS.

The following table shows how the disk types map to the effective disk type:

ONTAP disk type	Effective disk type
FCAL	SAS
SAS	SAS
ATA	FSAS
BSAS	FSAS
FCAL and SAS	SAS
MSATA	MSATA
FSAS	FSAS

What compatible spare disks are

In System Manager, compatible spare disks are disks that match the properties of other

disks in the aggregate. When you want to increase the size of an existing aggregate by adding HDDs (capacity disks) or change the RAID type of an aggregate from RAID4 to RAID-DP, the aggregate must contain sufficient compatible spare disks.

Disk properties that must match are the disk type, disk size (can be a higher size disk in case the same disk size is not available), disk RPM, checksum, node owner, pool, and shared disk properties. If you use higher sized disks, you must be aware that disk downsizing occurs and the size of all disks are reduced to the lowest disk size. Existing shared disks are matched with higher size non-shared disks, and the non-shared disks are converted to shared disks and added as spares.

If RAID mixing options, such as disk type mixing and disk RPM mixing, are enabled for the RAID group, the disk type and disk RPM of the existing disks of the aggregate are matched with the effective disk type and effective disk RPM of the spare disks to obtain compatible spares.

Related information

[Adding capacity disks](#)

[Editing aggregates](#)

How System Manager works with hot spares

A hot spare is a disk that is assigned to a storage system but not used by any RAID group. Hot spares do not contain any data and are assigned to a RAID group when a disk failure occurs in the RAID group. System Manager uses the largest disk as the hot spare.

When there are different disk types in the RAID group, the largest-sized disk of each disk type is left as the hot spare. For example, if there are 10 SATA disks and 10 SAS disks in the RAID group, the largest-sized SATA disk and the largest-sized SAS disk are serve as hot spares.

If the largest-sized disk is partitioned, then the hot spares are provided separately for partitioned and non-partitioned RAID groups. If the largest-sized disk is unpartitioned, then a single spare disk is provided.

The largest-sized non-partitioned disk is left as a hot spare if there are root partitions in the disk group. When a non-partitioned disk of the same size is not available, then spare root partitions are left as hot spares for the root partitioned group.

A single spare disk can serve as a hot spare for multiple RAID groups. System Manager calculates the hot spares based on the value set in the option `raid.min_spare_count` at the node level. For example, if there are 10 SSDs in an SSD RAID group and the option `raid.min_spare_count` is set to 1 at the node level, System Manager leaves 1 SSD as the hot spare and uses the other 9 SSDs for SSD-related operations. Similarly, if there are 10 HDDs in an HDD RAID group and the option `raid.min_spare_count` is set to 2 at the node level, System Manager leaves 2 HDDs as hot spares and uses the other 8 HDDs for HDD-related operations.

System Manager enforces the hot spare rule for RAID groups when you create an aggregate, edit an aggregate, and when you add HDDs or SSDs to an aggregate. The hot spare rule is also used when you create a storage pool or add disks to an existing storage pool.

There are exceptions to the hot spare rule in System Manager:

- For MSATA or disks in a multi-disk carrier, the number of hot spares is twice the value set at the node level and the number must not be less than 2 at any time.

- Hot spares are not used if the disks are part of array LUNs or virtual storage appliances.

Rules for displaying disk types and disk RPM

When you are creating an aggregate and adding capacity disks to an aggregate, you should understand the rules that apply when disk types and disk RPM are displayed.

When the disk type mixing and the disk RPM mixing options are not enabled, the actual disk type and actual disk RPM are displayed.

When these mixing options are enabled, the effective disk type and effective disk RPM are displayed instead of the actual disk type and actual disk RPM. For example, when the disk mixing option is enabled, System Manager displays BSAS disks as FSAS. Similarly, when the disk RPM mixing option is enabled, if the RPM of the disks is 10K and 15K, System Manager displays the effective RPM as 10K.



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.


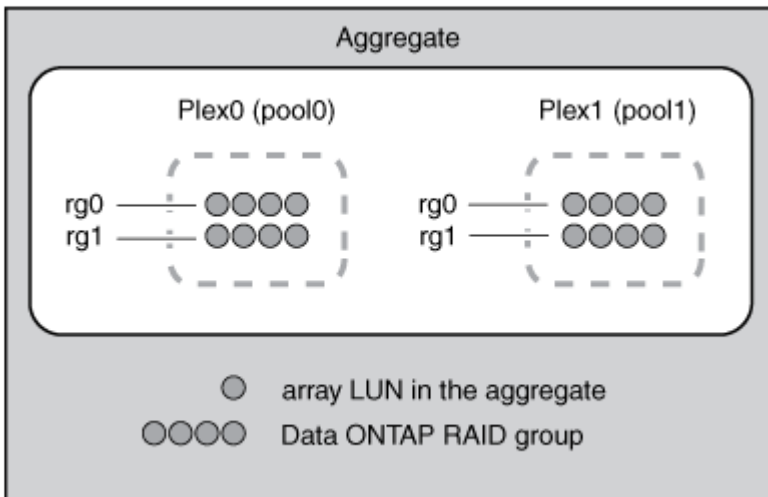
When a mirrored aggregate is created (or when a second plex is added to an existing unmirrored aggregate), ONTAP copies the data in the original plex (plex0) to the new plex (plex1). The plexes are physically separated (each plex has its own RAID groups and its own pool), and the plexes are updated simultaneously. 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.



What a FabricPool is

FabricPool is a hybrid storage solution that uses an all flash (all SSD) aggregate as the performance tier and an object store as the cloud tier. Data in a FabricPool is stored in a tier based on whether it is frequently accessed or not. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

Related information

[Adding a cloud tier](#)

Storage recommendations for creating aggregates

Starting with System Manager 9.4, you can create aggregates based on storage recommendations. However, you must determine whether creating aggregates based on storage recommendations is supported in your environment. If your environment does not support creating aggregates based on storage recommendations, you must decide the RAID policy and disk configuration, and then create the aggregates manually.

System Manager analyzes the available spare disks in the cluster and generates a recommendation about how the spare disks should be used to create aggregates according to best practices. System Manager displays the summary of recommended aggregates including their names and usable size.

In many cases, the storage recommendation will be optimal for your environment. However, if your cluster is running ONTAP 9.3 or earlier, or if your environment includes the following configurations, you must create aggregates manually:

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

In addition, if any of the following disk conditions exist in your environment, you must rectify the disk conditions before you use the storage recommendation to create aggregates:

- Missing disks
- Fluctuation in spare disk numbers
- Unassigned disks
- Non-zeroed spares (for ONTAP versions earlier than 9.6)
- Disks that are undergoing maintenance testing

Related information

[Disk and aggregate management](#)

[Zeroing spare disks](#)

Storage Tiers window

You can use the Storage Tiers window to view cluster-wide space details and to add and view aggregate details.

The Internal Tier panel, or the Performance Tier panel if the cluster has all flash (all SSD) aggregates, displays cluster-wide space details such as the sum of the total sizes of all of the aggregates, the space used by the

aggregates in the cluster, and the available space in the cluster.

The Cloud Tier panel displays the total licensed cloud tiers in the cluster, the licensed space that is used in the cluster, and the licensed space that is available in the cluster. The Cloud Tier panel also displays the unlicensed cloud capacity that is used.

Aggregates are grouped by type, and the aggregate panel displays details about the total aggregate space, space used, and the available space. If inactive (cold) data is available on a solid-state drive (SSD) or All Flash FAS aggregate, the amount of space it uses is also displayed. You can select the aggregate and perform any of the aggregate-related actions.

Command buttons

- **Add Aggregate**

Enables you to create an aggregate.

- **Actions**

Provides the following options:

- **Change status to**

Changes the status of the selected aggregate to one of the following statuses:

- **Online**

Read and write access to the volumes that are contained in this aggregate is allowed.

- **Offline**

Read and write access is not allowed.

- **Restrict**

Some operations such as parity reconstruction are allowed, but data access is not allowed.

- **Add Capacity**

Enables you to add capacity (HDDs or SSDs) to existing aggregates.

- **Add Cache**

Enables you to add cache disks (SSDs) to existing HDD aggregates or Flash Pool aggregates.

You cannot add cache disks to FabricPool-enabled aggregates.

This option is not available for a cluster containing nodes with All Flash Optimized personality.

- **Mirror**

Enables you to mirror the aggregates.

- **Volume Move**

Enables you to move a FlexVol volume.

Details area

You can click the aggregate name to view detailed information about the aggregate.

- **Overview tab**

Displays detailed information about the selected aggregate, and displays a pictorial representation of the space allocation of the aggregate, the space savings of the aggregate, and the performance of the aggregate.

- **Disk Information tab**

Displays the disk layout information for the selected aggregate.

- **Volumes tab**

Displays details about the total number of volumes on the aggregate, the total aggregate space, and the space committed to the aggregate.

- **Performance tab**

Displays graphs that show the performance metrics of the aggregates, including throughput and IOPS. Performance metrics data for read, write, and total transfers is displayed for throughput and IOPS, and the data for SSDs and HDDs is recorded separately.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. If you change the client time zone or the cluster time zone, you should refresh your browser to view the updated graphs.

Related information

[Adding a cloud tier](#)

[Attaching an aggregate to a cloud tier](#)

[Deleting a cloud tier](#)

[Editing a cloud tier](#)

[Provisioning storage through aggregates](#)

[Deleting aggregates](#)

[Editing aggregates](#)

Copyright Information

Copyright © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.