



Manage physical storage

System Manager Classic

NetApp
January 21, 2022

Table of Contents

- Managing physical storage 1
 - Storage tiers 1
 - Configure and managing cloud tiers 16
 - Aggregates 23
 - Storage pools 26
 - Disks 33
 - Array LUNs 40
 - Nodes 43
 - Hardware Cache 45
 - Events 46
 - System alerts 48
 - AutoSupport notifications 52
 - Jobs 55
 - Flash Pool statistics 57

Managing physical storage

You can use System Manager to manage physical storage such as aggregates, storage pools, disks, array LUNs, nodes, Flash Cache, events, system alerts, AutoSupport notifications, jobs, and Flash Pool statistics.

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](#)

[What compatible spare disks are](#)

[Storage Tiers 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

[Aggregates window](#)

[Storage Tiers window](#)

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	<ol style="list-style-type: none">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: <ol style="list-style-type: none">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.


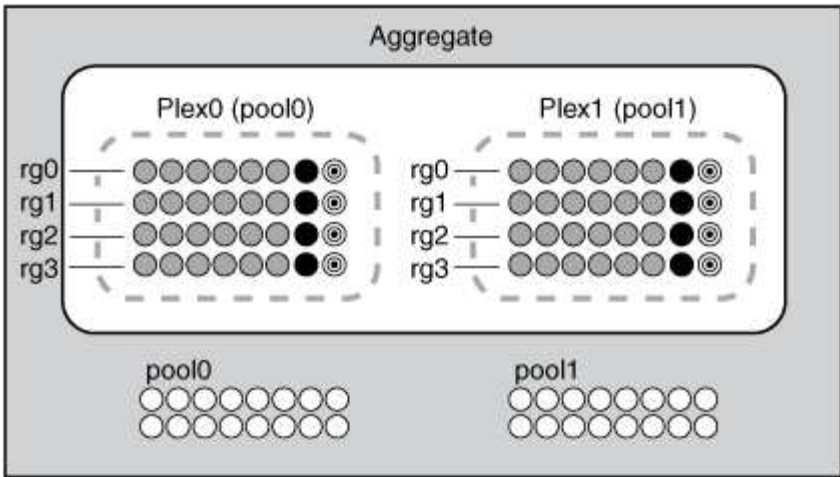
How mirrored aggregates work

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

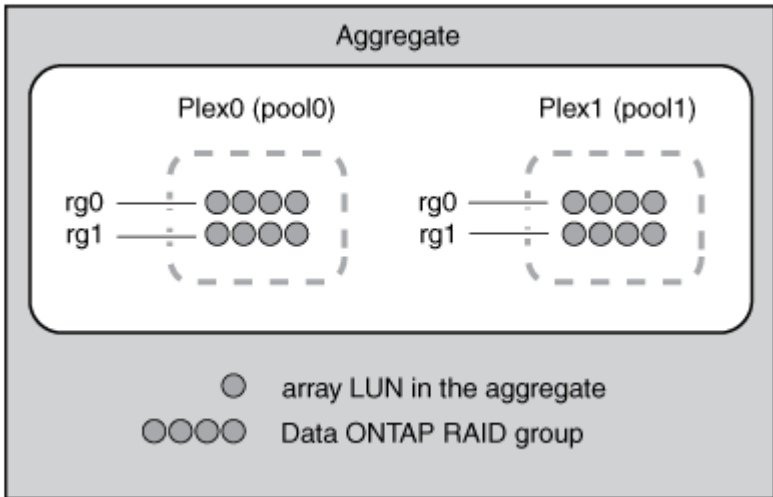
When a mirrored aggregate is created (or when a second plex is added to an existing unmirrored aggregate), ONTAP copies the data in the original plex (plex0) to the new plex (plex1). The plexes are physically separated (each plex has its own RAID groups and its own pool), and the plexes are updated simultaneously.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

[Attaching an aggregate to 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](#)


Configure and managing cloud tiers

Storing data in tiers can enhance the efficiency of your storage system. You manage storage tiers by using FabricPool-enabled aggregates. Cloud tiers store data in a tier based on whether the data is frequently accessed.

Before you begin

- You must be running ONTAP 9.2 or later.

- You must have all flash (all SSD) aggregates



Add a cloud tier

You can use System Manager to add a cloud tier to an SSD aggregate or a virtual machine disk (VMDK) aggregate. Cloud tiers provide storage for infrequently used data.

Before you begin

- You must have the access key ID and secret key to connect to the object store.
- You must have created a bucket inside the object store.
- Network connectivity must exist between the cluster and the cloud tier.
- If communication between the cloud tier and the cluster is encrypted using SSL or TLS, the required certificates must be installed.

About this task

The following object stores can be used as cloud tiers:

- StorageGRID
- Alibaba Cloud (Starting with System Manager 9.6)

- Amazon Web Services (AWS) Simple Storage Service (S3)
- Amazon Web Services (AWS) Commercial Cloud Service (C2S)
- Microsoft Azure Blob storage
- IBM Cloud
- Google Cloud



- Azure Stack, which is an on-premises Azure service, is not supported.
- If you want to use any object store as a cloud tier, other than StorageGRID, you must have the FabricPool capacity license. You can add the license by clicking **Add License**.
- If you want to use an IBM Cloud Object Storage environment (such as Cleversafe), with FabricPool, you should specify a certification authority (CA) certificate. You can specify the CA certificate by moving the **Object Store Certificate** toggle button and specifying the certificate credentials.

Steps

1. Click **Storage > Aggregates & Disks > Cloud Tier**.
2. Click **Add**.

A dialog box appears that lists all supported object store providers.

3. From the list, select the object store provider that you want to designate as the cloud tier.


The Add Cloud Tier window is displayed.

4. Select a region from the drop-down list in the **Region** field.

Based on your selection, the **Service Name (FQDN)** field is automatically populated with the server endpoint.

5. Specify the access key ID of the cloud tier, the secret key of the cloud tier, and the container name.

If you have selected AWS Commercial Cloud Service (C2S) as the type, you must specify the CAP URL, server CA certificates, and client certificates.

6. If you want to modify any of the following settings, then click the Advanced Options icon  to display the **Advanced Options** dialog window where you can make the changes:

- The port number used to access the cloud tier
- Enable or disable the **SSL** option that lets you transfer data securely to the cloud tier

7. If you want to add a cloud tier for StorageGRID or you want to use IBM Cloud Object Storage environment (such as Cleversafe) with FabricPool, you should specify a CA certificate. Specify the CA certificate by moving the **Object Store Certificate** toggle button and copying the contents of the certificate. Then paste the certificate contents in the signed certification.

8. From the **IPspace** list, select the IPspace that is used to connect to the cloud tier.

9. Click **Save** to save the cloud tier.

10. Click **Save and Attach Aggregates** to save the cloud tier and to attach aggregates to the cloud tier.

Related information

[What cloud tiers and tiering policies are](#)

[What a FabricPool is](#)

[Installing a CA certificate if you use StorageGRID](#)

[Storage Tiers window](#)

Attaching an aggregate to a cloud tier

You can use System Manager to attach an All Flash aggregate to a cloud tier. You can store infrequently used data in cloud tiers.

Before you begin

You must have added a cloud tier to the cluster.

Steps

1. Click **Storage > Aggregates & Disks > Cloud Tier**.
2. In the **Used in Aggregates** column, click **Attach Aggregates**.

The Attach Aggregates window is displayed.

3. Select the aggregate that you want to attach to the cloud tier.
4. Click **Save**.

Related information

[What cloud tiers and tiering policies are](#)

[What a FabricPool is](#)

[Storage Tiers window](#)

Provisioning storage by creating a FabricPool-enabled aggregate manually

You can use System Manager to create a FabricPool-enabled aggregate to attach a cloud tier to the SSD aggregate.

Before you begin

- You must have created a cloud tier and attached it to the cluster in which the SSD aggregate resides.
- An on-premises cloud tier must have been created.
- A dedicated network connection must exist between the cloud tier and the aggregate.

About this task

The following object stores can be used as cloud tiers:

- StorageGRID
- Alibaba Cloud (Starting with System Manager 9.6)
- Amazon Web Services (AWS) Simple Storage Service (S3)
- Amazon Web Services (AWS) Commercial Cloud Service (C2S)

- Microsoft Azure Blob storage
- IBM Cloud
- Google Cloud



- Azure Stack, which is an on-premises Azure services, is not supported.
- If you want to use any object store as a cloud tier, other than StorageGRID, you must have the FabricPool capacity license.

Steps

1. Create a FabricPool-enabled aggregate by using one of the following methods:
 - Click **Applications & Tiers > Storage Tiers > Add Aggregate**.
 - Click **Storage > Aggregate & Disks > Aggregates > Create**.
2. Enable the **Manually Create Aggregate** option to create an aggregate.
3. Create a FabricPool-enabled aggregate:
 - a. Specify the name of the aggregate, the disk type, and the number of disks or partitions to include in the aggregate.



Only all flash (all SSD) aggregates support FabricPool-enabled aggregates.

The minimum hot spare rule is applied to the disk group that has the largest disk size.

- b. Modify the RAID configuration of the aggregate:
 - i. Click **Change**.
 - ii. In the Change RAID Configuration dialog box, specify the RAID type and the RAID group size.

Shared disks support two RAID types: RAID-DP and RAID-TEC.
 - iii. Click **Save**.
4. Select the **FabricPool** checkbox, and then select a cloud tier from the list.
 5. Click **Create**.

Changing the tiering policy of a volume

You can use System Manager to change the default tiering policy of a volume to control whether the data of the volume is moved to the cloud tier when the data becomes inactive.

Steps

1. Click **Storage > Volumes**.
2. From the drop-down menu in the **SVM** field, select **All SVMs**.
3. Select the volume for which you want to change the tiering policy, and then click **More Actions > Change Tiering Policy**.
4. Select the required tiering policy from the **Tiering Policy** list, and then click **Save**.

Editing a cloud tier

You can use System Manager to modify the configuration information of cloud tier. The configuration details that you can edit include the name, fully qualified domain name (FQDN), port, access key ID, secret key, and object store certificate.

Steps

1. Click **Storage > Aggregates & Disks > Cloud Tier**.
2. Select the cloud tier that you want to edit, and then click **Edit**.
3. In the **Edit Cloud Tier** window, modify the cloud tier name, FQDN, port, access key ID, secret key, and object store certificate, as required.

If you have selected AWS Commercial Cloud Service (C2S) cloud tier, you can modify the server CA certificates, and client certificates.

4. Click **Save**.

Related information

[Storage Tiers window](#)

Deleting a cloud tier

You can use System Manager to delete a cloud tier that you no longer require.

Before you begin

You must have deleted the FabricPool-enabled aggregate that is associated with the cloud tier.

Steps

1. Click **Storage > Aggregates & Disks > Cloud Tier**.
2. Select the cloud tier that you want to delete, and then click **Delete**.

Related information

[Storage Tiers window](#)

What cloud tiers and tiering policies are

Cloud tiers provide storage for infrequently accessed data. You can attach an all-flash (all-SSD) aggregate to a cloud tier to store infrequently used data. You can use tiering policies to decide whether data should be moved to a cloud tier.

You can set one of the following tiering policies on a volume:

- **Snapshot-only**

Moves the Snapshot copies of only those volumes that are currently not being referenced by the active file system. Snapshot-only policy is the default tiering policy.

- **Auto**

Moves the inactive (cold) data and the Snapshot copies from the active file system to the cloud tier.

- **Backup (for System Manager 9.5)**

Moves the newly transferred data of a data protection (DP) volume to the cloud tier.

- **All (starting with System Manager 9.6)**

Moves all data to the cloud tier.

- **None**

Prevents the data on the volume from being moved to a cloud tier.

Related information

[Adding a cloud tier](#)

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

What inactive (cold) data is

Infrequently accessed data in a performance tier is known as inactive (cold) data. By default, data that is not accessed for a period of 31 days becomes inactive.

Inactive data is displayed at the aggregate level, cluster level, and volume level. The inactive data for an aggregate or a cluster is displayed only if inactive scanning is complete on that aggregate or cluster. By default, inactive data is displayed for FabricPool-enabled aggregates and SSD aggregates. Inactive data is not displayed for FlexGroups.

Cloud Tier window

You can use System Manager to add, edit, and delete cloud tiers and to view cloud tier details.

The Cloud Tier window displays the total number of 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 window also displays the unlicensed cloud capacity that is used.

Command buttons

- **Add**

Enables you to add a cloud tier.

- **Attach Aggregates**

Enables you to attach aggregates to a cloud tier.

- **Delete**

Enables you to delete a selected cloud tier.

- **Edit**

Enables you to modify the properties of a selected cloud tier.

Details area

You can view detailed information about cloud tiers such as the list of cloud tiers, the details of the object stores, the aggregates used, and the used capacity.

If you create a cloud tier other than Alibaba Cloud, Amazon AWS S3, AWS Commercial Cloud Service (C2S), Google Cloud, IBM Cloud, Microsoft Azure Blob storage, or StorageGRID by using the command-line interface (CLI), this cloud tier is displayed as Others in System Manager. You can then attach aggregates to this cloud tier.

Aggregates

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

Aggregates window

You can use the Aggregates window to create, display, and manage information about aggregates.

Command buttons

- **Create**

Opens the Create Aggregate dialog box, which enables you to create an aggregate.

- **Edit**

Opens the Edit Aggregate dialog box, which enables you to change the name of an aggregate or the level of RAID protection that you want to provide for the aggregate.

- **Delete**

Deletes the selected aggregate.



This button is disabled for the root aggregate.

- **More 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.

This button 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.

- **Attach Cloud Tier**

Enables you to attach a cloud tier to the aggregate.

- **Refresh**

Updates the information in the window.

Aggregate list

Displays the name and the space usage information for each aggregate.

- **Status**

Displays the status of the aggregate.

- **Name**

Displays the name of the aggregate.

- **Node**

Displays the name of the node to which the disks of the aggregate are assigned.

This field is available only at the cluster level.

- **Type**

Displays the type of the aggregate.

This field is not displayed for a cluster containing nodes with All Flash Optimized personality.

- **Used (%)**

Displays the percentage of space that is used in the aggregate.

- **Available Space**

Displays the available space in the aggregate.

- **Used Space**

Displays the amount of space that is used for data in the aggregate.

- **Total Space**

Displays the total space of the aggregate.

- **FabricPool**

Displays whether the selected aggregate is attached to a cloud tier.

- **Cloud Tier**

If the selected aggregate is attached to a cloud tier, it displays the name of the cloud tier.

- **Volume Count**

Displays the number of volumes that are associated with the aggregate.

- **Disk Count**

Displays the number of disks that are used to create the aggregate.

- **Flash Pool**

Displays the total cache size of the Flash Pool aggregate. A value of NA indicates that the aggregate is not a Flash Pool aggregate.

This field is not displayed for a cluster containing nodes with All Flash Optimized personality.

- **Mirrored**

Displays whether the aggregate is mirrored.

- **SnapLock Type**

Displays the SnapLock type of the aggregate.

Details area

Select an aggregate to view information about the selected aggregate. You can click Show More Details to view detailed information about the selected 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 in IOPS and total data transfers.

- **Disk Information tab**

Displays disk layout information such as the name of the disk, disk type, physical size, usable size, disk position, disk status, plex name, plex status, RAID group, RAID type, and storage pool (if any) for the selected aggregate. The disk port that is associated with the disk primary path and the disk name with the disk secondary path for a multipath configuration are also displayed.

- **Volumes tab**

Displays details about the total number of volumes on the aggregate, 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. You should refresh your browser to view the updated graphs.

Related information

[Provisioning storage through aggregates](#)

[Deleting aggregates](#)

[Editing aggregates](#)

Storage pools

You can use System Manager to create storage pools to enable SSDs to be shared by multiple Flash Pool aggregates.

Related information

[Disk and aggregate management](#)

Create a storage pool

A storage pool is a collection of SSDs (cache disks). You can use System Manager to combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares between an HA pair for allocation to two or more Flash Pool aggregates at the same time.

Before you begin

- Both nodes of the HA pair must be up and running in order to allocate SSDs and SSD spares through a storage pool.
- Storage pools must have a minimum of 3 SSDs.

- All SSDs in a storage pool must be owned by the same HA pair.

About this task

System Manager enforces the hot spare rule for SSD RAID groups when you use SSDs for adding disks to a storage pool. For example, if there are 10 SSDs in the 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.

You cannot use partitioned SSDs when creating a storage pool by using System Manager.

Steps

1. Click **Storage > Aggregates & Disks > Storage Pools**.
2. In the **Storage Pools** window, click **Create**.
3. In the **Create Storage Pool** dialog box, specify the name for the storage pool, disk size, and the number of disks.
4. Click **Create**.

Related information

[Storage Pools window](#)

Add disks to a storage pool

You can add SSDs to an existing storage pool and increase its cache size by using System Manager.

Before you begin

Both nodes of the HA pair must be up and running in order to allocate SSDs and SSD spares through a storage pool.

About this task

- The SSDs that you add to a storage pool are distributed proportionally among the aggregates using the storage pool cache and to the free space of the storage pool.
- System Manager enforces the hot spare rule for SSD RAID groups when you use SSDs for adding disks to a storage pool.

For example, if there are 10 SSDs in the 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.

- You cannot use partitioned SSDs when adding disks to a storage pool by using System Manager.

Steps

1. Click **Storage > Aggregates & Disks > Storage Pools**.
2. In the **Storage Pools** window, select the storage pool, and then click **Add Disks**.
3. In the **Add Disks** dialog box, specify the number of disks that you want to add.
4. Click **Next**.
5. In the **Summary** dialog box, review how the cache is distributed among various aggregates and the free space of the storage pool.

6. Click **Add**.

Related information

[Storage Pools window](#)

Deleting storage pools

You might want to delete a storage pool when the cache of the storage pool is not optimal or when it is no longer used by any aggregate or Flash Pool aggregate. You can delete a storage pool by using the Delete Storage Pool dialog box in System Manager.

Before you begin

The storage pool must not be used by any aggregate.

Steps

1. Click **Storage > Aggregates & Disks > Storage Pools**.
2. In the **Storage Pools** window, select the storage pool that you want to delete, and then click **Delete**.
3. In the **Delete Storage Pool** dialog box, click **Delete**.

Related information

[Storage Pools window](#)

Use SSD storage pools

To enable SSDs to be shared by multiple Flash Pool aggregates, you can add the SSDs to a *storage pool*. After you add an SSD to a storage pool, you can no longer manage the SSD as a stand-alone entity. You must use the storage pool to assign or allocate the storage that is provided by the SSD.

You can create storage pools for a specific high-availability (HA) pair. Then, you can add allocation units from that storage pool to one or more Flash Pool aggregates that are owned by the same HA pair. Just as disks must be owned by the same node that owns an aggregate before the disks can be allocated to it, storage pools can provide storage only to the Flash Pool aggregates that are owned by one of the nodes that owns the storage pool.

If you have to increase the amount of Flash Pool cache on your system, you can add more SSDs to a storage pool, up to the maximum RAID group size for the RAID type of the Flash Pool caches that are using the storage pool. When you add an SSD to an existing storage pool, you increase the size of the storage pool's allocation units, including any allocation units that are already allocated to a Flash Pool aggregate.

You can use only one spare SSD for a storage pool, so that if an SSD in that storage pool becomes unavailable, ONTAP can use the spare SSD to reconstruct the partitions of the malfunctioning SSD. You do not have to reserve any allocation units as spare capacity; ONTAP can use only a full, unpartitioned SSD as a spare for the SSDs in a storage pool.

After you add an SSD to a storage pool, you cannot remove the SSD, just as you cannot remove disks from an aggregate. If you want to use the SSDs in a storage pool as discrete drives again, you must destroy all of the Flash Pool aggregates to which the storage pool's allocation units have been allocated, and then destroy the storage pool.

Requirements and best practices for using SSD storage pools

Some technologies cannot be combined with Flash Pool aggregates that use SSD storage pools.

You cannot use the following technologies with Flash Pool aggregates that use SSD storage pools for their cache storage:

- MetroCluster
- SyncMirror functionality

Mirrored aggregates can coexist with Flash Pool aggregates that use storage pools; however, Flash Pool aggregates cannot be mirrored.

- Physical SSDs

Flash Pool aggregates can use SSD storage pools or physical SSDs, but not both.

SSD storage pools must conform to the following rules:

- SSD storage pools can contain only SSDs; HDDs cannot be added to an SSD storage pool.
- All of the SSDs in an SSD storage pool must be owned by the same high-availability (HA) pair.
- You cannot use SSDs that have been partitioned for root-data partitioning in a storage pool.

If you provide storage from a single storage pool to two caches with different RAID types, and you expand the size of the storage pool beyond the maximum RAID group size for RAID4, the extra partitions in the RAID4 allocation units go unused. Therefore, it is a best practice to keep your cache RAID types homogenous for a storage pool.

You cannot change the RAID type of cache RAID groups that are allocated from a storage pool. You set the RAID type for the cache before adding the first allocation units, and you cannot change the RAID type later.

When you create a storage pool or add SSDs to an existing storage pool, you must use the same size SSDs. If a failure occurs and no spare SSD of the correct size exists, ONTAP can use a larger SSD to replace the failed SSD. However, the larger SSD is right-sized to match the size of the other SSDs in the storage pool, resulting in lost SSD capacity.

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

Considerations for when to use SSD storage pools

SSD storage pools provide many benefits, but they also introduce some restrictions that you should be aware of when deciding whether to use SSD storage pools or dedicated SSDs.

SSD storage pools make sense only when they are providing cache to two or more Flash Pool aggregates. SSD storage pools provide the following benefits:

- Increased storage utilization for SSDs used in Flash Pool aggregates

SSD storage pools reduce the overall percentage of SSDs needed for parity by enabling you to share parity SSDs between two or more Flash Pool aggregates.

- Ability to share spares between HA partners

Because the storage pool is effectively owned by the HA pair, one spare, owned by one of the HA partners, can function as a spare for the entire SSD storage pool if needed.

- Better utilization of SSD performance

The high performance provided by SSDs can support access by both controllers in an HA pair.

These advantages must be weighed against the costs of using SSD storage pools, which include the following items:

- Reduced fault isolation

The loss of a single SSD affects all RAID groups that include one of its partitions. In this situation, every Flash Pool aggregate that has cache allocated from the SSD storage pool that contains the affected SSD has one or more RAID groups in reconstruction.

- Reduced performance isolation

If the Flash Pool cache is not properly sized, there can be contention for the cache between the Flash Pool aggregates that are sharing it. This risk can be mitigated with proper cache sizing and QoS controls.

- Decreased management flexibility

When you add storage to a storage pool, you increase the size of all Flash Pool caches that include one or more allocation units from that storage pool; you cannot determine how the extra capacity is distributed.

Considerations for adding SSDs to an existing storage pool versus creating a new one

You can increase the size of your SSD cache in two ways—by adding SSDs to an existing SSD storage pool or by creating a new SSD storage pool. The best method for you depends on your configuration and plans for the storage.

The choice between creating a new storage pool and adding storage capacity to an existing one is similar to deciding whether to create a new RAID group or add storage to an existing one:

- If you are adding a large number of SSDs, creating a new storage pool provides more flexibility because you can allocate the new storage pool differently from the existing one.
- If you are adding only a few SSDs, and increasing the RAID group size of your existing Flash Pool caches is not an issue, then adding SSDs to the existing storage pool keeps your spare and parity costs lower, and automatically allocates the new storage.

If your storage pool is providing allocation units to Flash Pool aggregates whose caches have different RAID types, and you expand the size of the storage pool beyond the maximum RAID4 RAID group size, the newly added partitions in the RAID4 allocation units are unused.

Why you add disks to storage pools

You can add SSDs to an existing storage pool and increase its cache size. When you add SSDs to a 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 of the storage pool.

If the allocation units of the storage pool are not yet allocated, adding SSDs to that storage pool does not affect the SSD cache size.

When you add SSDs to an existing 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.

How storage pool works

A *storage pool* is a collection of SSDs. You can combine SSDs to create a storage pool, which enables you to share the SSDs and SSD spares across multiple Flash Pool aggregates, at the same time.

Storage pools consist of allocation units, which you can use to provide SSDs and SSD spares to aggregates or to increase the existing SSD size.

After you add an SSD to a storage pool, you can no longer use the SSD as an individual disk. You must use the storage pool to assign or allocate the storage provided by the SSD.

Related information

[Provisioning storage by creating a Flash Pool aggregate manually](#)

[Provisioning cache by adding SSDs](#)

Storage Pools window

You can use the Storage Pools window to create, display, and manage a dedicated cache of SSDs, also known as *storage pools*. These storage pools can be associated with a non-root aggregate to provide SSD cache and with a Flash Pool aggregate to increase its size.

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

Command buttons

- **Create**

Opens the Create Storage Pool dialog box, which enables you to create a storage pool.

- **Add Disks**

Opens the Add Disks dialog box, which enables you to add cache disks to a storage pool.

- **Delete**

Deletes the selected storage pool.

- **Refresh**

Updates the information in the window.

Storage pools list

- **Name**

Displays the name of the storage pool.

- **Total Cache**

Displays the total cache size of the storage pool.

- **Spare Cache**

Displays the available spare cache size of the storage pool.

- **Used Cache (%)**

Displays the percentage of used cache size of the storage pool.

- **Allocation Unit**

Displays the minimum allocation unit of the total cache size that you can use to increase the size of your storage pool.

- **Owner**

Displays the name of the HA pair or the node with which the storage pool is associated.

- **State**

Displays the state of the storage pool, which can be Normal, Degraded, Creating, Deleting, Reassigning, or Growing.

- **Is Healthy**

Displays whether storage pool is healthy or not.

Details tab

Displays detailed information about the selected storage pool, such as the name, health, storage type, disk count, total cache, spare cache, used cache size (in percent), and allocation unit. The tab also displays the names of the aggregates that are provisioned by the storage pool.

Disks tab

Displays detailed information about the disks in the selected storage pool, such as the names, disk types, useable size, and total size.

Related information

[Adding disks to a storage pool](#)

[Creating a storage pool](#)

[Deleting storage pools](#)

Disks

You can use System Manager to manage disks.

Related information

[Disk and aggregate management](#)

[FlexArray virtualization installation requirements and reference](#)

[ONTAP concepts](#)

Reassigning disks to nodes

You can use System Manager to reassign the ownership of spare disks from one node to another node to increase the capacity of an aggregate or storage pool.

About this task

- You can reassign disks if the following conditions are true:
 - The container type of the selected disks must be “spare” or “shared”.
 - The disks must be connected to nodes in an HA configuration.
 - The disks must be visible to the node.
- You *cannot* reassign a disk if the following conditions are true:
 - The container type of the selected disk is “shared”, and the data partition is not spare.
 - The disk is associated with a storage pool.
- You cannot reassign the data partition of shared disks if storage failover is not enabled on the nodes that are associated with the shared disks.
- For partition disks, you can reassign only the data partition of the disks.
- For MetroCluster configurations, you cannot use System Manager to reassign disks.

You must use the command-line interface to reassign disks for MetroCluster configurations.

Steps

1. Click **Storage > Aggregates & Disks > Disks**.
2. In the **Disks** window, select the **Inventory** tab.
3. Select the disks that you want to reassign, and then click **Assign**.
4. In the **Warning** dialog box, click **Continue**.
5. In the **Assign Disks** dialog box, select the node to which you want to reassign the disks.
6. Click **Assign**.

Viewing disk information

You can use the Disks window in System Manager to view the name, size, and container details of disks along with graphical information about capacity disks and cache disks.

Steps

1. Click **Storage > Aggregates & Disks > Disks**.
2. Select the disk that you want to view information about from the displayed list of disks.
3. Review the disk details.

Related information

[Disks window](#)

How ONTAP reports disk types

ONTAP associates a type with every disk. ONTAP reports some disk types differently than the industry standards; you should understand how ONTAP disk types map to industry standards to avoid confusion.

When ONTAP documentation refers to a disk type, it is the type used by ONTAP unless otherwise specified. *RAID disk types* denote the role that a specific disk plays for RAID. RAID disk types are not related to ONTAP disk types.

For a specific configuration, the disk types that are supported depend on the storage system model, the shelf type, and the I/O modules that are installed in the system.

The following tables show how ONTAP disk types map to industry standard disk types for the SAS and FC storage connection types, and for storage arrays.

SAS-connected storage

ONTAP disk type	Disk class	Industry standard disk type	Description
BSAS	Capacity	SATA	Bridged SAS-SATA disks with added hardware to enable them to be plugged into a SAS-connected storage shelf
FSAS	Capacity	NL-SAS	Near Line SAS
MSATA	Capacity	SATA	SATA disk in multi-disk carrier storage shelf
SAS	Performance	SAS	Serial-Attached SCSI
SSD	Ultra-performance	SSD	Solid-state drives

FC-connected storage

ONTAP disk type	Disk class	Industry standard disk type
ATA	Capacity	SATA
FCAL	Performance	FC

Storage arrays

ONTAP disk type	Disk class	Industry standard disk type	Description
LUN	N/A	LUN	Logical storage device that is backed by storage arrays and used by ONTAP as a disk. These LUNs are referred to as <i>array LUNs</i> to distinguish them from the LUNs that ONTAP serves to clients.

Related information

[NetApp Hardware Universe](#)

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

Minimum number of hot spares required for disks

Having insufficient spares increases the risk of a disk failure with no available spare, resulting in a degraded RAID group. A spare disk is also required to provide important information (a *core file*) to technical support in case of a controller disruption.

MSATA disks, or disks in a multi-disk carrier, should have four hot spares during steady state operation, and you should never allow the number of MSATA hot spares to dip below two.

For RAID groups composed of SSDs, you should have at least one spare disk.

For all other ONTAP disk types, you should have at least one matching or appropriate hot spare available for each kind of disk installed in your storage system. However, having two available hot spares for all disks provides the best protection against disk failure. Having at least two available hot spares provides the following benefits:

- When you have two or more hot spares for a data disk, ONTAP can put that disk into the maintenance center if required.

ONTAP uses the maintenance center to test suspect disks and to take offline any disk that shows problems.

- Having two hot spares means that when a disk fails, you still have a spare disk available if another disk fails before you replace the first failed disk.

A single spare disk can serve as a hot spare for multiple RAID groups. However, if any disk in those RAID groups fails, then no spare disk is available for any future disk failures or for a core file until the spare disk is replaced. Therefore, it is a best practice to have more than one spare.

Spare requirements for multi-disk carrier disks

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

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

If two disks fail at the same time with only two available hot spares, ONTAP might not be able to swap the contents of both the failed disk and its carrier mate to the spare disks. This scenario is called a *stalemate*. If this happens, you are notified through EMS messages and AutoSupport messages. When the replacement carriers become available, you must follow the instructions that are provided by the EMS messages or you must contact technical support to recover from the stalemate.

Shelf configuration requirements for multi-disk carrier storage shelves

You can combine multi-disk carrier disk shelves with single-disk carrier disk shelves (standard disk shelves) on the same storage system and within in the same stack.

Determine when it is safe to remove a multi-disk carrier

Removing a multi-disk carrier before it is safe to do so can result in one or more RAID groups becoming degraded, or possibly even a storage disruption. System Manager enables you to determine when it is safe to remove a multi-disk carrier.

When a multi-disk carrier has to be replaced, the following events must have occurred before you can remove the carrier safely:

- An AutoSupport message must have been logged indicating that the carrier is ready to be removed.
- An EMS message must have been logged indicating that the carrier is ready to be removed.
- The state of both disks in the carrier must be displayed as `broken` in the Disks window.

You must remove the disks only after the carrier mate of a failed disk is evacuated. You can click Details to view the disk evacuation status in the Properties tab of the Disks window.

- The fault LED (amber) on the carrier must be lit continuously indicating that it is ready for removal.
- The activity LED (green) must be turned off indicating there is no disk activity.
- The shelf digital display only shows the shelf ID number.



You cannot reuse the carrier mate of a failed disk. When you remove a multi-disk carrier that contains a failed disk, you must replace it with a new carrier.

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 the same number of disks.

While you can have up to 50% less or more than the number of disks in different RAID groups on one aggregate, this might lead to performance bottlenecks in some cases, so is best avoided.

- The recommended range of RAID group disk numbers 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 disk numbers, you should choose the larger number of disks.

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.

Disks window

You can use the Disks window to view all the disks in your storage system.

Command buttons

- **Assign**

Assigns or reassigns the ownership of the disks to a node.

This button is enabled only if the container type of the selected disks is unassigned, spare, or shared.

- **Zero Spares**

Erases all the data, and formats the spare disks and array LUNs.

- **Refresh**

Updates the information in the window.

Tabs

Summary

Displays detailed information about the disks in the cluster, including the size of the spare disks and assigned disks. The tab also graphically displays information about spare disks, aggregates, and root aggregates for HDDs and information about spare disks, disks in a storage pool, aggregates, Flash Pool aggregates, and root aggregates for cache disks (SSDs).

The HDD panel is not displayed for systems with All Flash Optimized personality.

The details panel provides additional information about partitioned and unpartitioned spare disks (disk type, node, disk size, RPM, checksum, number of available disks, and spare capacity), in tabular format.

Inventory

- **Name**

Displays the name of the disk.

- **Container Type**

Displays the purpose for which the disk is used. The possible values are Aggregate, Broken, Foreign, Label Maintenance, Maintenance, Shared, Spare, Unassigned, Volume, Unknown, and Unsupported.

- **Partition Type**

Displays the partition type of the disk.

- **Node Name**

Displays the name of the node that contains the aggregate.

This field is available only at the cluster level.

- **Home owner**

Displays the name of the home node to which this disk is assigned.

- **Current owner**

Displays the name of the node that currently owns this disk.

- **Root owner**

Displays the name of the node that currently owns the root partition of this disk.

- **Data Owner**

Displays the name of the node that currently owns the data partition of this disk.

- **Data1 Owner**

Displays the name of the node that currently owns the data1 partition of the disk.

- **Data2 Owner**

Displays the name of the node that currently owns the data2 partition of the disk.

- **Storage Pool**

Displays the name of the storage pool with which the disk is associated.

- **Type**

Displays the type of the disk.

- **Firmware Version**

Displays the firmware version of the disk.

- **Model**

Displays the model of the disk.

- **RPM**

Displays the effective speed of the disk drive when the option `raid.mix.hdd.rpm.capacity` is enabled, and displays the actual speed of the disk drive when the option `raid.mix.hdd.rpm.capacity` is disabled.

This field is not applicable to SSDs.

- **Effective Size**

Displays the usable space available on the disk.

- **Physical Space**

Displays the total physical space of the disk.

- **Shelf**

Displays the shelf on which the physical disks are located.

This field is hidden by default.

- **Bay**

Displays the bay within the shelf for the physical disk.

This field is hidden by default.

- **Pool**

Displays the name of the pool to which the selected disk is assigned.

This field is hidden by default.

- **Checksum**

Displays the type of the checksum.

This field is hidden by default.

- **Carrier ID**

Specifies information about disks that are located within the specified multi-disk carrier. The ID is a 64-bit value.

This field is hidden by default.

Inventory details area

The area below the inventory tab displays detailed information about the selected disk, including information about the aggregate or volume (if applicable), vendor ID, zeroing state (in percent), serial number of the disk, and error details in case of a broken disk. For shared disks, the Inventory details area displays the names of all the aggregates, including the root and the non-root aggregates.

Related information

[Viewing disk information](#)

Array LUNs

You can use System Manager to assign array LUNs to an existing aggregate and manage array LUNs.

Related information

[FlexArray virtualization installation requirements and reference](#)

Assigning array LUNs

You can use System Manager to assign unassigned array LUNs to an existing aggregate to increase the size of the aggregate.

About this task

- You can assign array LUNs if the following conditions are true:
 - The container type of the selected array LUNs must be “unassigned”.
 - The disks must be connected to nodes in an HA pair.
 - The disks must be visible to the node.

- For MetroCluster configurations, you cannot use System Manager to assign array LUNs as spares.

You must use the command-line interface instead.

Steps

1. Click **Storage > Aggregates & Disks > Array LUNs**.
2. Select the array LUNs, and then click **Assign**.
3. In the **Assign Array LUNs** dialog box, select the node to which you want to assign the array LUNs.
4. Click **Assign**.

Reassigning spare array LUNs to nodes

You can use System Manager to reassign the ownership of spare array LUNs from one node to another to increase the capacity of an aggregate.

About this task

- You can reassign array LUNs if the following conditions are true:
 - The container type of the selected array LUNs must be “spare”.
 - The disks must be connected to nodes in an HA pair.
 - The disks must be visible to the node.
- For MetroCluster configurations, you cannot use System Manager to reassign array LUNs as spares.

You must use the command-line interface instead.

Steps

1. Click **Storage > Aggregates & Disks > Array LUNs**.
2. Select the spare array LUNs that you want to reassign, and then click **Assign**.
3. In the **Warning** dialog box, click **Continue**.
4. In the **Assign Array LUNs** dialog box, select the node to which you want to reassign the spare array LUNs.
5. Click **Assign**.

Zeroing spare array LUNs

You can use System Manager to erase all the data and to format the spare array LUNs by writing zeros to the array LUNs. These array LUNs can then be used in new aggregates.

About this task

When you zero the spare array LUNs, all the spares in the cluster, including disks, are zeroed. You can zero the spare array LUNs for a specific node or for the entire cluster.

Steps

1. Click **Storage > Aggregates & Disks > Array LUNs**.
2. Click **Zero Spares**.
3. In the **Zero Spares** dialog box, select a node or “All nodes” from which you want to zero the array LUNs.

4. Select the **Zero all non-zeroed spares** check box to confirm the zeroing operation.
5. Click **Zero Spares**.

About disks and array LUNs

A disk is the basic unit of storage for storage systems that use ONTAP software to access native disk shelves. An array LUN is the basic unit of storage that third-party storage arrays provide to storage systems that run ONTAP software.

ONTAP software enables you to assign ownership to your disks and array LUNs, and to add them to an aggregate. ONTAP software also provides a number of ways to manage your disks, including removing them, replacing them, and sanitizing them. Because array LUNs are provided by the third-party storage array, you use the third-party storage array for all other management tasks for array LUNs.

You can create an aggregate using either disks or array LUNs. After you have created the aggregate, you manage it using ONTAP software in exactly the same way, whether it was created from disks or array LUNs.

Array LUNs window

The Array LUNs window enables you to assign ownership to your array LUNs and to add them to an aggregate.

The Array LUNs link in the left navigation pane is displayed only if there are any spare array LUNs, or if the V_StorageAttach license is installed.

Command buttons

- **Assign**

Enables you to assign or reassign the ownership of array LUNs to a node.

- **Zero Spares**

Erases all the data, and formats the spare array LUNs and disks.

- **Refresh**

Updates the information in the window.

Array LUN list

Displays information such as the name, state, and vendor for each array LUN.

- **Name**

Specifies the name of the array LUN.

- **State**

Specifies the state of the array LUN.

- **Vendor**

Specifies the name of the vendor.

- **Used Space**

Specifies the space used by the array LUN.

- **Total Size**

Specifies the size of the array LUN.

- **Container**

Specifies the aggregate to which the array LUN belongs.

- **Node name**

Specifies the name of the node to which the array LUN belongs.

- **Home owner**

Displays the name of the home node to which the array LUN is assigned.

- **Current owner**

Displays the name of the node that currently owns the array LUN.

- **Array name**

Specifies the name of the array.

- **Pool**

Displays the name of the pool to which the selected array LUN is assigned.

Details area

The area below the Array LUNs list displays detailed information about the selected array LUN.

Nodes

You can use System Manager to view the details of the nodes in the cluster.

Initializing the ComplianceClock time

You can use System Manager to initialize the ComplianceClock time to the current cluster time. You must initialize the ComplianceClock time in order to create SnapLock aggregates.

Before you begin

The SnapLock license must be installed.

About this task

You cannot modify or stop the ComplianceClock time after it is initialized.

Steps

1. Click **Storage > Nodes**.
2. Select the node, and then click **Initialize ComplianceClock**.
3. In the **Initialize ComplianceClock** dialog box, click **Yes** to initialize the ComplianceClock time to the current cluster time.

Nodes window

You can use the Nodes window to view the details of the nodes in a cluster.

Command buttons

- **Initialize ComplianceClock**

Initializes the ComplianceClock of the selected node to the current value of the system clock.

- **Refresh**

Updates the information in the window.

Nodes list

- **Name**

Displays the name of the node.

- **State**

Displays the state of the node (whether the node is up or down).

- **Up Time**

Displays the duration for which the node is up.

- **ONTAP Version**

Displays the ONTAP version that is installed on the node.

- **Model**

Displays the platform model number of the node.

- **System ID**

Displays the ID of the node.

- **Serial No**

Displays the serial number of the node.

Details area

Displays detailed information about the selected node.

- **Details tab**

Displays information related to the selected node such as the name of the node, the state of the node, and the duration for which the node is up.

- **Performance tab**

Displays the throughput, IOPS, and latency of the selected node.

Changing the client time zone or the cluster time zone impacts the performance metrics graphs. You should refresh your browser to view the updated graphs.

Hardware Cache

You can use System Manager to manage Hardware Cache modules.



Flash Cache is known as Hardware Cache in System Manager.

Enabling or disabling Flash Cache modules

You can enable or disable the external cache functionality for a storage system that has a Flash Cache module installed by using System Manager. You can enable Flash Cache modules based on the workload requirements of your storage system.

Steps

1. Click **Configuration > Hardware Cache**
2. If you want to modify a module or modules, move the slider button to enable or disable each module, as required.

How Flash Cache modules work

Using Flash Cache modules improves the performance of a storage system. The impact of using Flash Cache modules is displayed on the Hardware Cache window.

You can configure Flash Cache modules and disks based on the workload requirements of a storage system. By determining the read workload (number of read operations) served by Flash Cache modules and disks, you can analyze the performance of the storage system.

Flash Cache modules do not contain any data during the storage system boot or when control is returned to the storage system after a takeover event. Therefore, disks serve all the data read requests of the storage system.

The Flash Cache module is slowly populated with data when data read requests are served. Because the data read requests served by Flash Cache modules are faster than those served by the disks, the performance of the storage system improves.

Data read requests served by the Flash Cache module replace the data read requests served by the disks,

and therefore, the performance improvement in the storage system is directly related to the disk reads that are replaced. To understand the impact of Flash Cache modules on storage system performance, you must view the read workload graph in the Hardware Cache window when the Flash Cache module contains data.

Hardware Cache window

You can use the Hardware Cache window to enable or disable Flash Cache modules for a storage system that has a Flash Cache module installed. You can also view the read-workload statistics.

Module Information

- **Storage system name**

The name of the storage system that has a Flash Cache module installed displays under the graphic.

- **Enable/Disable toggle button**

Move the toggle button to enable or disable the module.

- **Size**

The size of the module in gigabytes. If there are multiple Flash Cache module cards, the total cache size from all of the cards is displayed.



The Flash Cache module size that is displayed differs from the actual size for the following reasons: - System Manager reports only the usable capacity that is provided by ONTAP. - A portion of the total capacity is reserved for storing metadata.

- **Model Names**

The model names of the modules.

- **System Read Latency**

Displays the average read latency in milliseconds.

Cache Read Workload

Indicates storage system performance by displaying a graph specifying the rate of the read workload that is served by the disks and the Flash Cache module.

Events

You can use System Manager to view the event log and event notifications.

Events window

You can use the Events window to view the event log and event notifications.

Command buttons

- **Refresh**

Updates the information in the window.

Events list

- **Time**

Displays the time when the event occurred.

- **Node**

Displays the node and the cluster on which the event occurred.

- **Severity**

Displays the severity of the event. The possible severity levels are:

- Emergency

Specifies that the event source unexpectedly stopped, and the system experienced unrecoverable data loss. You must take corrective action immediately to avoid extended downtime.

- Alert

Specifies that the event source has an alert, and action must be taken to avoid downtime.

- Critical

Specifies that the event source is critical, and might lead to service disruption if corrective action is not taken immediately.

- Error

Specifies that the event source is still performing, and a corrective action is required to avoid service disruption.

- Warning

Specifies that the event source experienced an occurrence that you must be aware of. Events of this severity might not cause service disruption; however, corrective action might be required.

- Notice

Specifies that the event source is normal, but the severity is a significant condition that you must be aware of.

- Informational

Specifies that the event source has an occurrence that you must be aware of. No corrective action might be required.

- Debug

Specifies that the event source includes a debugging message.

By default, the alert severity type, emergency severity type, and the error severity type are displayed.

- **Source**

Displays the source of the event.

- **Event**

Displays the description of the event.

Details area

Displays the event details, including the event description, message name, sequence number, message description, and corrective action for the selected event.

System alerts

You can use System Manager to monitor different parts of a cluster.

Related information

[System administration](#)

Acknowledging system health alerts

You can use System Manager to acknowledge and respond to system health alerts for subsystems. You can use the information displayed to take the recommended action and correct the problem reported by the alert.

Steps

1. Click **Events & Jobs > System Alerts**.
2. In the **System Alerts** window, click the arrow icon next to the name of subsystem.
3. Select the alert that you want to acknowledge, and then click **Acknowledge**.
4. Type your name, and then click **Acknowledge**.

Related information

[System Alerts window](#)

Suppressing system health alerts

You can use System Manager to suppress system health alerts that do not require any intervention from you.

Steps

1. Click **Events & Jobs > System Alerts**.
2. In the **System Alerts** window, click the arrow icon next to the name of subsystem.

3. Select the alert that you want to suppress, and then click **Suppress**.
4. Type your name, and then click **Suppress**.

Related information

[System Alerts window](#)

Deleting system health alerts

You can use System Manager to delete system health alerts to which you have already responded.

Steps


1. Click **Events & Jobs > System Alerts**.
2. In the **System Alerts** window, click the arrow icon next to the name of subsystem.
3. Select the alert that you want to delete, and then click **Delete**.
4. Click **OK**.

Related information

[System Alerts window](#)

Available cluster health monitors

There are several health monitors that monitor different parts of a cluster. Health monitors help you to recover from errors within ONTAP systems by detecting events, sending alerts to you, and deleting events as they clear.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
Cluster switch(cluster-switch)	Switch (Switch-Health)	<p>Monitors cluster network switches and management network switches for temperature, utilization, interface configuration, redundancy (cluster network switches only), and fan and power supply operation. The cluster switch health monitor communicates with switches through SNMP. SNMPv2c is the default setting.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Beginning with ONTAP 9.2, this monitor can detect and report when a cluster switch has rebooted since the last polling period.</p> </div>

MetroCluster Fabric	Switch	Monitors the MetroCluster configuration back-end fabric topology and detects misconfigurations such as incorrect cabling and zoning, and ISL failures.
MetroCluster Health	Interconnect, RAID, and storage	Monitors FC-VI adapters, FC initiator adapters, left-behind aggregates and disks, and inter-cluster ports
Node connectivity(node-connect)	CIFS nondisruptive operations (CIFS-NDO)	Monitors SMB connections for nondisruptive operations to Hyper-V applications.
Storage (SAS-connect)	Monitors shelves, disks, and adapters at the node level for appropriate paths and connections.	System
not applicable	Aggregates information from other health monitors.	System connectivity (system-connect)

Ways to respond to system health alerts

When a system health alert occurs, you can acknowledge it, learn more about it, repair the underlying condition, and prevent it from occurring again.

When a health monitor raises an alert, you can respond in any of the following ways:

- Get information about the alert, which includes the affected resource, alert severity, probable cause, possible effect, and corrective actions.
- Get detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already.
- Get health-related information about the state of the affected resource or subsystem, such as a specific shelf or disk.
- Acknowledge the alert to indicate that someone is working on the problem, and identify yourself as the “Acknowledger.”
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- Delete the alert, if the system did not automatically clear it.
- Suppress an alert to prevent it from affecting the health status of a subsystem.

Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur, but the subsystem health displays as “ok-with-suppressed.” when the suppressed alert occurs.

System Alerts window

You can use the System Alerts window to learn more about system health alerts. You can also acknowledge, delete, and suppress alerts from the window.

Command buttons

- **Acknowledge**

Enables you to acknowledge the selected alert to indicate that the problem is being addressed and identifies the person who clicks the button as the “Acknowledger.”

- **Suppress**

Enables you to suppress the selected alert to prevent the system from notifying you about the same alert again and identifies you as the “Suppressor.”

- **Delete**

Deletes the selected alert.

- **Refresh**

Updates the information in the window.

Alerts list

- **SubSystem (No. of Alerts)**

Displays the name of the subsystem, such as the SAS connection, switch health, CIFS NDO, or MetroCluster, for which the alert is generated.

- **Alert ID**

Displays the alert ID.

- **Node**

Displays the name of the node for which the alert is generated.

- **Severity**

Displays the severity of the alert as Unknown, Other, Information, Degraded, Minor, Major, Critical, or Fatal.

- **Resource**

Displays the resource that generated the alert, such as a specific shelf or disk.

- **Time**

Displays the time when the alert was generated.

Details area

The details area displays detailed information about the alert, such as the time when the alert was generated and whether the alert has been acknowledged. The area also includes information about the probable cause and possible effect of the condition generated by the alert, and the recommended actions to correct the problem reported by the alert.

Related information

[Acknowledging system health alerts](#)

[Suppressing system health alerts](#)

[Deleting system health alerts](#)


AutoSupport notifications

You can use System Manager to configure AutoSupport notifications that help you to monitor your storage system health.

Setting up AutoSupport notifications

You can use the Edit AutoSupport Settings dialog box in System Manager to set up AutoSupport notifications by specifying an email address from which email notifications are sent and adding multiple email host names.

Steps

1. Click  > **AutoSupport**.
2. Select the node, and then click **Edit**.
3. In the **Email Recipient** tab, type the email address from which email notifications are sent, specify the email recipients and the message content for each email recipient, and add the mail hosts.

You can add up to five email addresses of the host names.

4. In the **Others** tab, select a transport protocol for delivering the email messages from the drop-down list and specify the HTTP or HTTPS proxy server details.
5. Click **OK**.
6. Verify that configuration you have set for AutoSupport notification is set up correctly in the **AutoSupport** window.


Enabling or disabling AutoSupport settings

You can enable or disable AutoSupport settings on your storage system by using System Manager. AutoSupport messages enable you to monitor your storage system health or send notifications to technical support and your internal support organization.

About this task

The AutoSupport option is enabled by default.


Steps

1. Click  > **AutoSupport**.
2. Select the node, and then click **Enable** or **Disable**.
3. Click **OK**.
4. Verify that the AutoSupport status correctly displays the change you made.

Add AutoSupport email recipients

You can use the **Email Recipient** tab of the Edit AutoSupport Settings dialog box in System Manager to add email addresses of the recipients of AutoSupport notifications.


Steps

1. Click  > **AutoSupport**.
2. Select the node, and then click **Edit**.
3. In the **Email Recipient** tab, type the address of the email recipient, specify whether the recipient receives a full message or a short message, and then click **Add**.
4. Click **OK**.
5. Verify that the details you specified are displayed in the **AutoSupport** window.

Testing AutoSupport settings

You can use the AutoSupport Test dialog box in System Manager to test that you have configured the AutoSupport settings correctly.

Steps

1. Click  > **AutoSupport**.
2. Select the node, and then click **Test**.
3. In the **AutoSupport Test** dialog box, enter the AutoSupport subject text “Test AutoSupport” or any text that notifies the recipients that you are testing the AutoSupport settings.
4. Click **Test**.

An email message with the subject “Test AutoSupport” or the text that you typed in the **AutoSupport subject** field is sent to the specified recipients.

Generating AutoSupport data

You can use System Manager to generate AutoSupport data for a single node or multiple nodes to monitor their health and to send notifications to technical support.

Steps

1. Click  > **AutoSupport**.
2. Select the node, and then click **AutoSupport Request** > **Generate AutoSupport**.

By default, the AutoSupport data is generated for all nodes.

3. In the **Generate AutoSupport** dialog box, perform the following steps:
 - a. If you want to generate AutoSupport data for a specific node, clear the **Generate Autosupport data**

for all nodes check box, and then select the node.

b. Type the case number.

4. Click **Generate**.

5. In the **Confirmation** dialog box, click **OK**.

Viewing AutoSupport summary

System Manager enables you to view the status and details of all the previous AutoSupport data in order to review the data that has been sent to technical support. You can also view the information to understand the health and performance of your storage system.

Steps

1. Click  > **AutoSupport**.

2. Select the node, and then click **AutoSupport Request** > **View Previous Summary**.

The AutoSupport data for all the nodes is displayed.

3. Click **OK**.

AutoSupport severity types

AutoSupport messages have severity types that help you understand the purpose of each message—for example, to draw immediate attention to an emergency problem, or only to provide information.

Messages have one of the following severities:

- **Alert:** Alert messages indicate that a next-higher level event might occur if you do not take some action.
You must take an action against alert messages within 24 hours.
- **Emergency:** Emergency messages are displayed when a disruption has occurred.
You must take an action against emergency messages immediately.
- **Error:** Error conditions indicate what might happen if you ignore.
- **Notice:** Normal but significant condition.
- **Info:** Informational message provides details about the issue, which you can ignore.
- **Debug:** Debug-level messages provide instructions you should perform.

If your internal support organization receives AutoSupport messages through email, the severity appears in the subject line of the email message.

AutoSupport window

The AutoSupport window enables you to view the current AutoSupport settings for your system. You can also change your system's AutoSupport settings.

Command buttons

- **Enable**

Enables AutoSupport notification. **Enable** is the default.

- **Disable**

Disables AutoSupport notification.

- **Edit**

Opens the Edit AutoSupport Settings dialog box, which enables you to specify an email address from which email notifications are sent and to add multiple email addresses of the host names.

- **Test**

Opens the AutoSupport Test dialog box, which enables you to generate an AutoSupport test message.

- **AutoSupport Request**

Provides the following AutoSupport requests:

- **Generate AutoSupport**

Generates AutoSupport data for a selected node or all nodes.

- **View Previous Summary**

Displays the status and details of all the previous AutoSupport data.

- **Refresh**

Updates the information in the window.

Details area

The details area displays AutoSupport setting information such as the node name, AutoSupport status, transport protocol used, and name of the proxy server.

Jobs

You can use System Manager to manage job tasks such as displaying job information and monitoring the progress of a job.

Jobs

Jobs are asynchronous task and typically long-running volume operations, such as copying, moving, or mirroring data. Jobs are placed in a job queue and are run when resources are available. The cluster administrator can perform all the tasks related to job management.

A job can be one of the following categories:

- A *server-affiliated* job is placed in queue by the management framework to be run in a specific node.
- A *cluster-affiliated* job is placed in queue by the management framework to be run in any node in the cluster.
- A *private* job is specific to a node and does not use the replicated database (RDB) or any other cluster mechanism.

You require the advanced privilege level or higher to run the commands to manage private jobs.

You can manage jobs in the following ways:

- Displaying job information, including the following:
 - Jobs on a per-node basis
 - Cluster-affiliated jobs
 - Completed jobs
 - Job history
- Monitoring a job's progress
- Displaying information about the initialization state for job managers.

You can determine the outcome of a completed job by checking the event log.

Job window

You can use the Job window to manage job tasks such as displaying job information and monitoring the progress of a job.

Command button

- **Refresh**

Updates the information in the window.

Tabs

- **Current Jobs**

This tab displays information about the job tasks that are in progress.

- **Job History**

This tab displays information about all the jobs.

Job list

- **Job ID**

Displays the ID of the job.

- **Start Time**

Displays the start time of the job.

- **Job Name**

Displays the name of the job.

- **Node**

Displays the name of the node.

- **State**

Displays the state of the job.

- **Job Description**

Displays the description of the job.

- **Progress**

Displays the state of the job.

- **Schedule Name**

Displays the name of the schedule.

Flash Pool statistics

You can use System Manager to view the real-time SSD tier read and write workloads for a selected Flash Pool aggregate.

Flash Pool aggregate Statistics window

You can view the real-time SSD tier read and write workloads for a selected Flash Pool aggregate.

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

- **Displaying Statistics for Flash Pool aggregate**

From the list of Flash Pool aggregates, you can select the Flash Pool aggregate whose statistics you want to view.

SSD Cache Read Workload

Displays a graphical view of the total read requests that are sent to the Flash Pool aggregate in comparison with the read operations that are performed by the SSD tier.

SSD Cache Write Workload

Displays a graphical view of the total write requests that are sent to the Flash Pool aggregate in comparison with the write operations that are performed by the SSD tier.

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.