

Dell Unity™ Family Configuring Pools

Version 5.x

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Additional resources.....	6
Chapter 1: Introduction to pools.....	7
About pools.....	7
Dynamic pools.....	7
Traditional pools.....	8
Storage tiers.....	9
Pool best practices.....	10
Spare drive policy (physical deployments only).....	11
Considerations for expanding pools.....	12
Chapter 2: Configuring pools.....	13
Create a dynamic pool.....	13
Create a traditional pool in physical deployments.....	14
Create a pool in virtual deployments.....	14
Enable automatic snapshot deletion for a pool.....	15
Monitor pool space allocation and usage.....	15
View drive information	15
Change pool properties.....	16
Expand a pool in physical deployments.....	16
Expand a pool in virtual deployments.....	17
Delete a pool.....	17
Chapter 3: Introduction to Data Efficiency.....	18
Optimizing drive performance using the FAST Cache and FAST VP (supported physical deployments only).....	18
Automatic snapshot deletion.....	19
About Data Reduction and Advanced Deduplication.....	19
Chapter 4: Managing the FAST Cache (supported physical deployments only).....	21
About the FAST Cache	21
FAST Cache best practices.....	22
Create the FAST Cache	22
Enable or disable the FAST Cache for a pool	22
View system-level FAST Cache information	22
View pool-level Fast Cache Information	23
Expand or shrink the FAST Cache capacity	23
Delete the FAST Cache	23
Chapter 5: Managing FAST VP (supported physical deployments only).....	24
About FAST VP	24
Data relocation with FAST VP	24
FAST VP tiering policy	25
Enable or disable scheduled data relocations at the system level.....	26

Stop a pool from participating in scheduled data relocations	26
Manually initiate or stop data relocations for a pool	26
Pause or resume data relocations for all pools	27
Change system-level data relocation information	27
Change data relocation information for a pool	27
View pool-level FAST VP information	28
Chapter 6: Understanding RAID (physical deployments only)	29
About RAID	29
RAID configurations	30
RAID for dynamic pools in Unisphere	31
Drive IOPS by RAID type	31
Chapter 7: Configuring pools using the CLI.....	33
Configure pools automatically	33
Initiate automatic pool configuration.....	34
View configuration settings for automatic pool creation.....	34
Configure custom pools.....	35
Create pools.....	38
Change pool settings.....	41
Add drives to pools.....	43
View pools.....	44
Delete pools.....	47
Manage pool tiers.....	48
View storage tiers.....	49
View pool resources.....	50
View storage profiles (physical deployments only).....	52
Manage drive groups (physical deployments only).....	54
View drive groups.....	55
View recommended drive group configurations.....	56
Manage storage system capacity settings.....	57
View system capacity settings.....	58
Manage system tier capacity settings.....	58
View system tier capacity.....	59
Change disk settings (virtual deployments only).....	59
Pool use cases.....	60
Create a pool using drives with specific characteristics.....	60
Configure a dynamic pool.....	61
Configure a traditional pool for an all-Flash model.....	62
Add drives to an existing pool.....	63
The show action command.....	64
Chapter 8: Managing the FAST Cache using the CLI (supported physical deployments only).....	66
Manage FAST Cache.....	66
Create FAST Cache.....	66
View FAST Cache settings.....	67
Extend FAST Cache.....	67
Shrink FAST Cache.....	68
Delete FAST Cache.....	68

Manage FAST Cache storage objects (physical deployments only).....	69
View FAST Cache storage objects.....	69
Chapter 9: Managing FAST VP using the CLI (supported physical deployments only).....	71
Manage FAST VP pool settings.....	71
View FAST VP pool settings.....	72
Change FAST VP pool settings.....	73
Start data relocation.....	74
Stop data relocation.....	75
Manage FAST VP general settings.....	75
Change FAST VP general settings.....	76
View FAST VP general settings.....	77

As part of an improvement effort, revisions of the software and hardware are periodically released. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features. Contact your technical support professional if a product does not function properly or does not function as described in this document.

Where to get help

Support, product, and licensing information can be obtained as described below.

Product information

For product and feature documentation or release notes, go to Unity Technical Documentation at: dell.com/unitydocs.

Troubleshooting

For information about products, software updates, licensing, and service, go to Support (registration required) at: dell.com/support. After logging in, locate the appropriate product page.

Introduction to pools

Topics:

- [About pools](#)
- [Dynamic pools](#)
- [Traditional pools](#)
- [Storage tiers](#)
- [Pool best practices](#)
- [Spare drive policy \(physical deployments only\)](#)
- [Considerations for expanding pools](#)

About pools

A pool is a set of drives that provide specific storage characteristics for the resources that use them. For example, the pool configuration defines the types and capacities of the drives in the pool. For physical deployments, the pool configuration also defines the RAID configurations (RAID types and stripe widths) for these drives.

You choose which pool to use when you create a new storage resource.

NOTE: Before you create storage resources, you must configure at least one pool. You cannot shrink a pool, and you cannot change its storage characteristics without deleting the storage resources configured in the pool and the pool itself. However, you can add drives to expand the pool.

Pools generally provide optimized storage for a particular set of applications or conditions. When you create a storage resource for hosts to use, you must choose a pool with which to associate the storage resource. The storage that the storage resource uses is drawn from the specified pool. If there are multiple drive types on the system, you can define multiple tiers for the pool. In physical deployments, each tier can be associated with a different RAID type.

Unity supports two types of pools, depending on your Unity model.

- [Dynamic pools](#)
- [Traditional pools](#)

All-Flash models running Unity OE version 4.2 and later supports both dynamic and traditional pools. In these models, all new pools created in the Unisphere GUI are dynamic pools. New pools created using the Unisphere CLI or REST API can be dynamic pools (the default) or traditional pools. Although you cannot create traditional pools in the Unisphere GUI, you can use the GUI to manage existing traditional pools.

All-Flash models running Unity OE version 4.1x and earlier, and all hybrid and VSA models support traditional pools only.

Dynamic pools

In Unity All-Flash models running OE version 4.2.x or later, all new pools created in the Unisphere GUI are dynamic pools, and new pools created in the Unisphere CLI and REST API are dynamic pools by default. Dynamic pools implement advanced RAID technology. In dynamic pools, a RAID group is spread across drive extents in multiple drives. The required spare space is also spread across drive extents in multiple drives. When a drive fails, the extents on the failed drive are rebuilt to spare space extents within the pool.

NOTE: For Unity All-Flash models running OE version 4.2.x or later, you can create traditional pools using the Unisphere CLI or REST API.

Dynamic pools have the following advantages over traditional pools:

- Drives are not wasted, because there are no fixed spares. All drives in the system can be added to a pool. This prolongs the life of the drives in the pool, since the load is spread across additional drives.

- Rebuild times are usually much faster than with traditional pools. Since spare capacity for a dynamic pool is spread across multiple drives rather than concentrated in on a single hot spare drive, more drives contribute to the rebuilding process when a drive fails.
- Pools can usually be expanded based on desired capacity. For example, you can add one drive at a time to a dynamic pool, providing provisioning flexibility and cost savings.

In a dynamic pool, drives are grouped. When the number of drives reaches 64, a new group is created. Different groups follow the same common rule. In each group, you can generally provision a dynamic pool with any number of drives, as long as the minimum drive number is satisfied for each specified drive type/capacity combination. The minimum drive number is the selected stripe width for a drive type/capacity combination plus additional spare space count (for example, 6 drives for RAID 5 (4 + 1) by default (1 of 32), but 7 drives if spare space count is 2 of 32).

The following considerations apply to dynamic pools:

- Once a dynamic pool is created, you cannot change its RAID type or stripe width. However, if you expand the pool using a different drive type, the added drives can have a different stripe width.
- You cannot shrink a dynamic pool or change its storage characteristics without deleting the storage resources configured in the pool and the pool itself. However, you can add drives to expand the pool.
- You can mix Flash drives of the same drive type with different capacities when you provision a dynamic pool. However, if you do this, the system might not use the entire capacity of the larger drives. This depends on how many drives of each capacity are in the pool. The unused capacity in a dynamic pool might become available during a future pool expansion.

One drive's worth of capacity equal to that of the largest drive in the pool is set aside as spare space for every set of 32 drives in a dynamic storage pool by default. If using two drive's worth of capacity for every 32 drives, the spare space will be doubled. For example, by default a dynamic pool with 1 to 32 drives of a given drive type has 1 drive's worth of spare space, while a dynamic pool with 33 to 64 drives of a given drive type has 2 drive's worth of spare space. However, if using 2 drive's worth of spare space, a dynamic pool with 2 to 32 drives of a given drive type has 2 drive's worth of spare space, while a dynamic pool with 33 to 64 drives of a given drive type has 4 drive's worth of spare space.

At a minimum, there must be the equivalent capacity of one spare drive per pool. Therefore, to minimize the amount of spare capacity required, it is recommended that you create dynamic pools with larger, rather than smaller, numbers of drives of the same drive type.

Dynamic pools can be homogeneous or heterogeneous. All drives in a homogeneous pool have the same drive type, such as SAS Flash 2 or SAS Flash 3 drives. Drives in a heterogeneous pool can include SAS Flash 2 drives, SAS Flash 3 drives, and SAS Flash 4 drives.

Traditional pools

Pools created in UnityVSA models, hybrid models, and Unity All-Flash models running OE version 4.1.x or earlier are traditional pools. For Unity All-Flash models running OE version 4.2.x or later, you can create traditional pools using the Unisphere CLI or REST API, but not the Unisphere GUI.

Traditional pools can be homogeneous or heterogeneous. All drives in a homogeneous pool have the same drive type, such as SAS drives or SAS Flash 2 drives. Drives in a heterogeneous pool have a mixture of drive types, such as a mixture of NL-SAS, SAS, and SAS Flash 2 drives. Traditional pools can also be All-Flash or hybrid. A hybrid pool contains a mixture of Flash and non-Flash drives. All supported drive types can be included in a hybrid pool, except for SAS Flash 4 drives, which must be in an All-Flash pool.

In physical deployments, storage in traditional pools is managed in RAID group units, where:

- A drive is consumed by a single RAID group.
- A RAID group is limited to a maximum of 16 drives and is composed of drives of the same type.
- Each tier supports a single RAID type.

Since storage in traditional pools is managed in RAID group units, adding capacity to a pool requires that you add drives in RAID group increments. For example, to add drives to a pool with RAID 5 (4+1), you must add at least 5 drives to the pool. As drive capacity increases, the minimum amount of storage that can be added to a pool and the cost of that storage become increasingly large.

The following considerations apply to traditional pools:

- Once a tier in a traditional pool is created, you cannot change the RAID type or stripe width of the existing drives in the tier. However, if you expand a tier within a traditional pool, you can specify a different stripe width for the newly-added drives. When you add a new tier to a traditional pool, you can specify a different RAID type, stripe width, or both for the newly-added drives.
- You cannot shrink a traditional pool or change its storage characteristics without deleting the storage resources configured in the pool and the pool itself. However, you can add drives to expand the pool.

With traditional pools, the storage system uses dedicated hot spares to replace a drive that has failed or faulted. Any unused drive in the system with the appropriate drive technology and size can be used to replace a failed or faulted drive in a pool. If a spare drive with the same type and size is not available, the system can use a larger drive of the same type. Because spare drives are dedicated hot spares, they cannot be used to improve pool performance or mitigate Flash drive wear. Also, when a drive fails or is faulted, the whole drive must be rebuilt on the spare drive. Therefore, the rebuild time can be very long, because it is limited by the performance of the single drive whose contents is being rebuilt. This can impact performance. It can also increase the chances of encountering additional drive failures during the rebuild process, which can lead to data loss.

Storage tiers

The storage tiers available for both physical and virtual deployments are described in the table below.

- For physical deployments, the storage tier is associated with the physical drive type.
- For virtual deployments, the storage tier is associated with the virtual disk's underlying characteristics and must be manually assigned.
- For both types of deployments, if FAST VP is installed on the system, you can create tiered pools to optimize drive utilization. A tiered pool consists of multiple drive types, such as SAS Flash 2 drives and SAS drives.

NOTE: For dynamic storage pools, **Unused Drives** does not include HDD system drives. For traditional storage pools, **Unused Drives** includes HDD system drives.

During creation of a dynamic pool, you can configure the hot spare capacity for each storage tier as follows:

- 1/32—Reserve 1 drive capacity for each 32 drives as a spare
- 2/32—Reserve 2 drives capacity for each 32 drives as spares

Once configured, the hot spare capacity configuration is permanent for the life of the pool, and cannot be changed.

The minimum number of drives to use for a dynamic pool is the selected RAID width plus one or two, depending on the spare policy used.

Table 1. Storage tier descriptions

Storage tier	Drive types	Description	Default RAID configuration (physical deployments only)
Extreme Performance tier	Solid state extreme performance drives. The following types are supported: <ul style="list-style-type: none"> • SAS Flash 2 • SAS Flash 3 • SAS Flash 4 	Provides very fast access times for resources subject to variable workloads. For example, databases can achieve their best performance when using SAS Flash drives. SAS Flash drives are more expensive than SAS drives per GB of storage. SAS Flash 2, SAS Flash 3, and SAS Flash 4 drives can be used together in the Extreme Performance tier. Only SAS Flash 2 drives can be used in the FAST Cache, and only SAS Flash 2 and SAS Flash 3 drives can be used in FAST VP.	RAID 5 (8 + 1).
Performance tier	SAS - Rotating performance drive	Provides high, all-around performance with consistent response times, high throughput, and good bandwidth at a mid-level price point. Performance tier storage is appropriate for database resources accessed centrally through a network.	RAID 5 (4 + 1).

Table 1. Storage tier descriptions (continued)

Storage tier	Drive types	Description	Default RAID configuration (physical deployments only)
Capacity tier	NL-SAS - Rotating capacity drive	<p>Provides the highest storage capacity with generally lower performance. Capacity storage is appropriate for storing large amounts of data that is primarily static (such as video files, audio files, and images) for users and applications that do not have strict performance requirements.</p> <p>NOTE: For data that changes or is accessed frequently, capacity tier storage has significantly lower performance.</p>	RAID 6 (6 + 2).

Pool best practices

Using fewer pools reduces complexity and increases flexibility. However, it may be appropriate to configure multiple pools, in order to:

- Separate workloads with different I/O profiles.
- Dedicate resources to meet specific performance goals.
- Separate resources for multi-tenancy.
- Create smaller failure domains.
- (Traditional hybrid pools only) Separate pools where FAST Cache is and is not active.

Pools must maintain free capacity in order to operate properly. By default, the storage system raises an alert if a pool has less than 30% free capacity, and will begin to automatically invalidate snapshots and replication sessions if a pool has less than 5% free capacity. It is recommended that a pool always have at least 10% free capacity.

All-Flash pools

All-Flash pools provide the highest level of performance in Unity. Use an all-Flash pool for applications that require the highest storage performance at the lowest response time. Note the following about all-Flash pools:

- Compression is only supported in all-Flash pools.
- Snapshots and replication operate most efficiently in all-Flash pools.
- FAST Cache and FAST VP are not applicable to all-Flash pools.

Follow these best practices for dynamic pools:

- Because of the way spare space is reserved for dynamic pools, it is recommended that you create dynamic pools with larger, rather than smaller, numbers of drives of the same drive type. Following this practice minimizes the amount of spare capacity required for a dynamic pool.
- Expanding a dynamic pool by adding multiples of the pool's RAID stripe width plus one allows space to be available faster than if you expand the pool by adding fewer drives. You can see the RAID stripe width for a pool with a single drive type by using the CLI.

For traditional pools, it is recommended that you use a single drive size and a single RAID width within a traditional all-Flash pool.

Hybrid pools

Hybrid pools typically provide greater capacity at a lower cost than all-Flash pools, but also typically have lower overall performance and higher response times. Use hybrid pools for applications that do not require consistently low response times, or that have large amounts of mostly inactive data.

It is recommended that you provision a Flash tier in hybrid pools. The Flash tier helps enable pool performance efficiencies, and improves response times when using data reduction and snapshots or replication, or both. In order to enable data reduction, the minimum recommended Flash capacity is at least 10% of the pool capacity. Otherwise, 5% of Flash is the minimum recommended Flash capacity for improving hybrid performance.

You can improve the performance of a hybrid pool by increasing the amount of capacity in the Flash tier, so that more of the active dataset resides on and is serviced by the Flash drives.

Hybrid pools can have up to three tiers (Extreme Performance, Performance, and Capacity). It is recommended that you use a single drive speed, size, and RAID width within a tier of a hybrid pool.

Spare drive policy (physical deployments only)

Traditional Pools

For traditional pools, any unused drive in the system with the appropriate drive technology and size, or larger, can be used to replace a failed or faulted drive. Most of the drive configurations require the use of hot sparing, except for the following:

- If the system has only 6 drives in total, and they are of the same type, you can configure RAID 10 (3+3) with no hot spare.
- If the system has only 8 drives in total, and they are of the same type, you can configure RAID 6 (6+2) with no hot spare.
- If the system has only 12 drives in total, and they are of the same type, you can configure RAID 6 (10+2) with no hot spare.

NOTE: When you expand traditional pools for a drive configuration, the system prevents you from configuring all available drives, allowing you to leave some drives as spares.

The spare drive policy follows these rules to determine how many drives are left as spares:

- For traditional pools, in general, the system reserves 1 spare drive for every group of 1 to 31 drives that have the same type, capacity, and rotational speed (or Flash type). For example, if there are 40 300-GB, 15K-RPM SAS drives in the system, the system reserves 2 of those drives as spares.
- The system does not reserve a spare drive for:
 - The FAST Cache.
 - A system drive, unless it has user data on it.
- Any unused non-system drive can become a spare drive.
- A system drive that does not contain user data can be a spare drive for a failed system drive that has user data.
- A spare drive can be used to replace a failed or faulted drive in the FAST Cache.
- When a spare drive swaps into a pool, it becomes a permanent member of that pool and cannot be used in another pool.
- When a broken drive is fixed or replaced, it can be a candidate for a spare drive or used in another pool.

Refer to the compatibility and interoperability documentation on the support website for a listing of basic platform and component support for the storage system, including capacity limits.

Dynamic Pools

For dynamic pools, the hot spare policy follows these rules to determine how much extent pool space is used as dynamic pool spare space:

- A dynamic pool is created by using a disk partner group (DPG). One DPG contains up to 64 drives. When DPG1 is full, DPG2 will be created. Drives in a DPG should have the same type, capacity, and rotational speed (or Flash type).
- In general, the system reserves 1 drive capacity spare space in a DPG with 1 to 32 drives, and 2 drives capacity spare space in a DPG of 33 to 64 drives. You also have the option of using 2 drive capacity spare spaces for a dynamic pool. Once configured, the hot spare capacity configuration is permanent for the life of the dynamic pool, and cannot be changed.
- A dynamic pool can consume all available drives, and does not reserve any hot spare drives outside of the pool.
- Compared to a traditional pool, dynamic pools use spare space inside the pool, instead of a hot spare drive.
- Hot spare space can be used only in dynamic pools with an Extreme Performance tier containing All-Flash drives, because by default All-Flash drives have mapped RAID enabled.

Considerations for expanding pools

Dynamic pools


You can usually expand dynamic pools based on desired capacity instead of by adding drives that are multiples of the pool's RAID stripe width plus one. For example, you can add two drives to a RAID 5 pool instead of six drives. The following exceptions apply:

- If you add a new drive type when you expand a pool, you must expand the pool with the minimum amount of drives required for the selected RAID configuration.
- Depending on the RAID configuration and number of drives in the pool, there are certain internal thresholds at which you must expand the pool with multiples of the pool's RAID stripe width plus one. The system tells you how many drives you need to add when you try to expand a pool that has reached one of these thresholds.

Expanding a pool by adding multiples of the pool's RAID strip width plus one allows space to be available faster than if you expand the pool by adding fewer drives. You can see the RAID stripe width for a pool by accessing the **RAID** tab on the pool's properties page.

Traditional pools

You expand a traditional pool by adding drives to the pool's existing tiers, adding new tiers that have available drives, or both. When you add drives, you must add them in multiples of the selected RAID stripe width.

 **NOTE:** You cannot expand an All-Flash pool in a hybrid model by adding SAS or NL-SAS drives if the pool contains:

- SAS Flash 4 drives.

 **NOTE:**

- Prior to Unity OE version 5.2.x, resources with data reduction enabled or were previously enabled and then disabled could only be expanded with additional flash drives.
- Starting with Unity OE version 5.2.x, expanding and converting an all-flash pool to a hybrid pool is supported. Data reduction-enabled resources, with or without advanced deduplication enabled, will no longer block pool expansions as long as the flash capacity is not less than 10% of the pool.
- By expanding a pool, if the additional capacity drops the flash capacity below 10%, a warning is displayed, and you cannot proceed.

Configuring pools

Topics:

- [Create a dynamic pool](#)
- [Create a traditional pool in physical deployments](#)
- [Create a pool in virtual deployments](#)
- [Enable automatic snapshot deletion for a pool](#)
- [Monitor pool space allocation and usage](#)
- [View drive information](#)
- [Change pool properties](#)
- [Expand a pool in physical deployments](#)
- [Expand a pool in virtual deployments](#)
- [Delete a pool](#)

Create a dynamic pool

Prerequisites

In Unity All-Flash models running OE version 4.2.x or later and Unity HFA models running OE version 5.2.x or later, all new pools that are created in the Unisphere UI, Unisphere CLI, and REST API are dynamic pools by default.


When you create a dynamic pool with NL-SAS drives, only RAID 6 is allowed when using the Unisphere UI. You can create a dynamic pool with other RAID types by using the UEMCLI for NL-SAS.

The minimum number of drives to use for a dynamic pool is the selected RAID width plus one drive.

Before you create a dynamic pool:

- If you have a hybrid model, find out whether the storage system is licensed for FAST Cache. To do this, select the Settings icon, and then select **Software and Licenses > License Information**. If the storage system is licensed for FAST Cache, it appears in the License Management table and you can choose whether to use it for the pool.
- If you have a hybrid model, find out whether the storage system is licensed for FAST VP. To do this, select the Settings icon, and then select **Software and Licenses > License Information**. If the storage system is licensed for FAST VP, it appears in the License Management table and you can choose whether to use it for the pool.
- Decide what RAID protection to use for the pool.
- Decide how many hot spare drives you want to use for the pool. The default is 1/32 (reserve 1 drive capacity for every 32 drives as a spare). You can also choose 2/32 (reserve 2 drives capacity for every 32 drives as spares).
- Decide whether to create a capability profile that has capabilities which are based on the pool configuration. To use the capability profile for VMware vVols, you must assign specific usage tags, which are propagated to the VMware vSphere environment, and can be used in policy profiles. The virtualization administrator and storage administrator should work together to define these tags.


About this task

 **NOTE:** Once a dynamic pool is created, you cannot change either the RAID type or hot spare configuration. Also, you cannot shrink a dynamic pool or change its storage characteristics without deleting the storage resources that are configured in the pool and the pool itself. However, you can add drives to expand the pool.

Steps

1. Under **Storage**, select **Pools**.
2. Select the **Add** icon.
3. Follow the steps in the wizard.

The number and types of drives you can choose is based on the RAID configuration, which you can optionally update by clicking **Change**. You can also set or change the hot spare capacity as part of changing the RAID configuration.

 **NOTE:** For dynamic storage pools, the **Unused Drives** value in the **Tiers** tab does not include HDD system drives. For traditional storage pools, the **Unused Drives** value in the **Tiers** tab includes HDD system drives.

Create a traditional pool in physical deployments

Prerequisites

All pools created with All-Flash models running Unity OE version 4.1.x or earlier and hybrid models running Unity OE version 5.1.x or earlier, support traditional pools only. With newer All-Flash models, you can create a traditional pool using the Unisphere CLI and REST API.

Before you create a traditional pool:

- If you have a hybrid model, find out whether the storage system is licensed for FAST Cache. To do this, select the Settings icon, and then select **Storage Configuration > FAST Cache**. If the storage system is licensed for FAST Cache, you can choose whether to use it for the pool.
- If you have a hybrid model, find out whether the storage system is licensed for FAST VP. To do this, select the Settings icon, and then select **Storage Configuration > FAST VP**. If the storage system is licensed for FAST VP you can choose which storage tiers to add to the pool.
- Decide whether to change the suggested RAID type of the tiers.
- Decide whether to create a capability profile that has capabilities based on the pool configuration. To use the capability profile for VMware vVols, you must assign specific usage tags, which are propagated to the VMware vSphere environment, and can be used in policy profiles. The virtualization administrator and storage administrator should work together to define these tags.

Steps

1. Under **Storage**, select **Pools**.
2. Select the **Add** icon.
3. Follow the steps in the wizard, taking into account the following considerations:
 - On the **Tiers** screen, you can only select multiple storage tiers if the system is a hybrid model that is licensed to use FAST VP. The wizard displays a maximum usable capacity for each selected tier, which it calculates based on the default RAID configuration. You can optionally change the RAID configuration for all selected tiers.
 - The number and types of drives you can choose is based on the RAID configuration.

Create a pool in virtual deployments


Prerequisites

All pools created in virtual deployments are traditional pools.

Before you create a traditional pool:

- From the storage administrator, obtain information about the underlying characteristics of the drives to use in the pool. You will use this information to assign tiers to the virtual disks that do not already have them assigned. You can create a multi-tier pool if the system has multiple underlying drive types.
- Decide whether to create a capability profile for VMware vVols that has capabilities based on the pool configuration. To use the capability profile, you must assign specific usage tags, which are propagated to the VMware vSphere environment, and can be used in policy profiles. The virtualization administrator and storage administrator should work together to define these tags.

About this task

 **NOTE:** You cannot shrink a pool or change its storage characteristics without deleting the storage resources configured in the pool and the pool itself. However, you can add drives to expand the pool.

Steps

1. Under **Storage**, select **Pools**.
2. Select the **Add** icon.
3. Select the tiers and virtual disks to use in the pool. Each virtual disk must have an assigned tier. If the virtual disks you want to include in the pool do not have assigned tiers, you must assign one. The tier you select for a virtual disk must be based on the underlying drive characteristics.
4. Optionally create a VMware capability profile for use by vVols, and specify usage tags for that profile.

Enable automatic snapshot deletion for a pool

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and click the **Edit** icon.
3. Select the **Snapshot Settings** tab.
4. In the **Automatically delete oldest snapshots** section, select whether you want the automatic deletion to occur when the total pool consumption space reaches a threshold you set, snapshot pool consumption space reaches a threshold you set, or both. For each option, you can also set the threshold when the automatic snapshot deletion stops.

Monitor pool space allocation and usage

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and then select the **Edit** icon.
3. Select **Usage > Capacity** to view pool space usage with metrics on: current pool capacity, space used, alert threshold, data reduction savings, and a graph showing the pool's used capacity over time.

Note the following:

- The **Non-base Space** field displays the space used by all snapshots and thin clones in the pool.
 - The **Used** field displays the percentage of the pool's total space that has been reserved by its associated storage resources. This total includes the space used by thin clones and snapshots. When this value is over 100%, the pool is oversubscribed.
 - The **Preallocated Space** displays the amount of remaining space in the pool that is reserved for, but not actively being used by, a storage resource.
 - The **Alert Threshold** field displays the percentage of storage allocation at which Unisphere generates notifications about the amount of space remaining in the pool. You can set this value between 50% and 84%.
4. Select **Storage Resources** to view a list of storage resources in the pool, along with the following metrics: total amount of pool space used by each resource and amount of pool space used for each resource's snapshots and thin clones.

View drive information

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > Drives** to view the information about the drives in the storage system:
 - For physical deployments this table includes the:
 - Drive type
 - Total number of drives
 - Number of unconfigured drives that are available for provisioning, excluding traditional pools. For dynamic storage pools, this value also excludes HDD system drives.
 - Number of unconfigured drives that are reserved as hot spares for traditional pools
 - Number of configured drives that are provisioned in dynamic pools
 - Number of configured drives that are provisioned in FAST Cache
 - Number of configured drives that are provisioned in traditional pools

- For virtual deployments, this table includes the size and WWN of the virtual drives that are available for use in a pool. Drives that are already used in a pool or that are faulted are not shown in the list.
2. To view information about unconfigured drives for a specific drive type and capacity, select the number in the **Unconfigured Available** or **Unconfigured Traditional Spare** column. The system displays the health status, name, type, drive size, and rotational speed (for hard drives) of the available drives.

Change pool properties

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and then select the **Edit** icon.
3. On the **General** tab, optionally change the pool name and description.
4. On the **Drives** tab, review the characteristics of the drives in the pool.
5. On the **Usage** tab, select **Capacity** to:
 - View the amount of free space in the pool.
 - View the subscribed pool capacity.
 - View the pool used capacity history.
 - View the space saved (if data reduction is enabled on a storage resource which supports it).
 - Change the alert threshold.

Note the following:

- The **Used** field displays the percentage of the pool's total space that its associated storage resources reserves. This total includes the space that thin clones and snapshots use. When this value is over 100%, the pool is oversubscribed.
 - The **Non-base Space** field displays the space that all snapshots and thin clones use in the pool.
 - The **Preallocated Space** displays the amount of remaining space in the pool that is reserved for a storage resource, but that the storage resource is not using.
 - The **Alert Threshold** field displays the percentage of storage allocation at which Unisphere generates notifications about the amount of space remaining in the pool. You can set this value between 50% and 84%.
6. On the **Usage** tab, select **Storage Resources** to view a list of storage resources in the pool, along with applicable metrics.
 7. On the **FAST VP** tab (which appears if FAST VP is licensed), view data relocation and tier information for the pool. You can also exclude the pool from data relocations, manually start or stop data relocations for the pool, and manually move data to different tiers. **Click Manage FAST VP** settings to manage system-wide FAST VP settings.
 8. On the **Snapshot Settings** tab, review and optionally change the properties for snapshot automatic deletion.
 9. On the **RAID** tab (dynamic pools only), view the:
 - Drive type for each drive in the pool.
 - RAID type and stripe width of each drive type.
 - Number of drives for each drive type in the pool.
 - Configured hot spare capacity.


Expand a pool in physical deployments


About this task

Expand a dynamic or traditional pool by adding drives to the pool tiers.


Steps

1. Under **Storage**, select **Pools**.
2. Select the pool to expand, and then select **Expand Pool**.

 **NOTE:** If the **Expand Pool** button is disabled, a tool tip will display with an explanation.
3. On the **Storage Tiers** tab, select the storage tiers for the drives you want to add to the pool.

 **NOTE:** For dynamic storage pools, **Unused Drives** does not include HDD system drives. For traditional storage pools, **Unused Drives** includes HDD system drives.

4. On the **Drives** tab, select the number of drives to add to each tier. You can select Flash drives by capacity.

 **NOTE:**

- The minimum number of drives to use for a dynamic pool is the selected RAID width plus one drive.
- When expanding a hybrid pool, and data reduction and advanced deduplication are enabled, or were previously enabled and then disabled, for a storage resource in the pool, the percentage of Flash drives in the pool must be 10% or higher. Otherwise, an error will occur.

Expand a pool in virtual deployments

About this task

Expand a pool by adding virtual disks to the pool. Before you expand a pool, obtain information from your storage administrator about the underlying characteristics of the virtual disks to use in the pool.

Steps

1. Under **Storage**, select **Pools**.
2. Select the pool to expand, and then select **Expand Pool**.
3. On the **Assign Tier to the Virtual Disk** page, assign or optionally change the storage tiers to which virtual disks are assigned. To do this, select the virtual disk, select the **Edit** icon, and then select the tier.
4. On the **Tiers** page, select one or more tiers to add to the pool.
5. On the **Virtual Disks** page, select the virtual disks to add to the pool.

Delete a pool

About this task

You can delete pools if they contain no allocated storage resources, including consistency groups, LUNs, file systems, VMware datastores, or NAS servers.

Steps

1. Under **Storage**, select **Pools**.
2. Select the pool to delete, and click the **Delete** icon.

Introduction to Data Efficiency

Topics:

- [Optimizing drive performance using the FAST Cache and FAST VP \(supported physical deployments only\)](#)
- [Automatic snapshot deletion](#)
- [About Data Reduction and Advanced Deduplication](#)

Optimizing drive performance using the FAST Cache and FAST VP (supported physical deployments only)

The FAST (Fully Automated Storage Tiering) Suite includes features that enable you to:

- Leverage SAS Flash 2 drives as additional read/write cache for improved performance (FAST Cache).
- Dynamically tier data across different types of drives (FAST VP). FAST VP can use all supported drive types except SAS Flash 4.

Comparison of FAST Cache and FAST VP

The following table describes the differences between the FAST Cache and FAST VP features:

Table 2. Differences between the FAST Cache and FAST VP features

FAST Cache	FAST VP
Enables SAS Flash 2 drives to be used as an additional read/write cache for the storage system.	Leverages pools to provide sub-LUN and file system tiering, which moves data to the appropriate tier based on the FAST VP tiering policy.
Caches 64-KB data chunks for higher performance.	Relocates 256-MB data chunks based on the FAST VP tiering policy and I/O activity.
Copies data from Hard Disk Drives (HDDs) to Flash drives when the data is accessed frequently.	Moves data between different storage tiers based on the FAST VP tiering policy.
Adapts continuously to changes in workload.	Uses a data relocation process to periodically make storage tiering adjustments. The data relocation schedule is configurable, and the data relocation can take place once a day or on an ongoing basis throughout the day.

Interoperability considerations

You can use FAST Cache and FAST VP functionality together to yield high performance and improve Total Cost of Ownership (TCO) for the storage system. It is recommended that you:

1. Use available SAS Flash 2 drives for the FAST Cache first, because this can benefit all storage resources in the storage system.
2. Supplement performance as needed by adding additional SAS Flash 2 drives to pool tiers for use by FAST VP.

For example, in scenarios where limited SAS Flash 2 drives are available, you can use SAS Flash 2 drives to create the FAST Cache, and you can apply FAST VP on a one- or two-tier pool (SAS and NL-SAS). From a performance point of view, FAST Cache dynamically provides performance benefits to bursts of data, while FAST VP moves "hotter" data to performance drives and "colder" data to capacity drives. From a TCO perspective, FAST Cache, with a small number of Flash drives, serves the data that is accessed most frequently, while FAST VP optimizes drive utilization and efficiency.

The FAST Cache feature is storage-tier-aware and works with FAST VP to make sure that the storage system resources are not wasted by unnecessarily copying data to FAST Cache, if it is already on a Flash drive. If FAST VP moves a chunk of data to the Extreme Performance Tier on a pool, the system will not copy that chunk of data into FAST Cache, even if FAST Cache criteria is met for promotion. This ensures that the storage system resources are not wasted in copying data from one Flash drive to another.

Automatic snapshot deletion

Automatic snapshot deletion is a space management feature used to automatically manage the number of snapshots in a pool. This feature is triggered when the total pool consumption or pool consumption by the snapshots reach a high threshold you define. The system automatically starts deleting old and expired snapshots until the pool space reaches a set threshold.

Expired snapshots are deleted first. If deleting the expired snapshots does not still result in reaching the thresholds, the system starts deleting the detached oldest snapshots with the automatic deletion option enabled. Automatic deletion does not apply to snapshots attached to hosts, including attached groups of snapshots on a consistency group. It also does not apply to system snapshots used for replication.

You can set a snapshot of a storage resource for automatic deletion when the snapshot expires (by setting a snapshot expiration date) or the associated pool reaches the auto delete threshold settings. The following table explains the automatic snapshot deletion options:


 **NOTE:** If a snapshot has automatic snapshot deletion enabled, it cannot be used as a source snapshot for creating or refreshing a thin clone.

Table 3. Automatic snapshot deletion options

Options	Description
Pool	Applies to snapshots of storage resources and pools. When the pool reaches a certain threshold, the system automatically deletes the snapshot.
Expiration time	Applies to snapshots of storage resources. This is the time when the snapshot expires.

About Data Reduction and Advanced Deduplication

A single switch called Data Reduction controls compression and deduplication of the blocks at an individual LUN and file system level. Adding a subsequent switch called Advanced Deduplication controls a deeper, dynamic ability for deduplicating data blocks within a given storage resource that do not contain internally-defined data patterns.

Applying data reduction saves space on the system. Data reduction includes savings from deduplicating blocks which contain internally-defined data patterns, including zero detection and compression. Data reduction is available for thin LUNs, thin file systems, and thin NFS and VMFS datastores. The thin file system must be created on a Unity system that is running OE version 4.2.x or later.

Data reduction is applied on all new incoming writes. Data that cannot be deduplicated or compressed is detected, and is written uncompressed.

Data that already exists on a LUN will not have data reduction applied after data reduction is enabled. In order to apply data reduction to existing LUN data, you must use the LUN Move feature to move the LUN's data to a destination LUN that has data reduction enabled.

The Advanced Deduplication option helps increase the storage efficiency of systems above what is achieved with data reduction on a LUN, file system, or datastore. It provides the ability to reduce the amount of storage needed for user data by keeping only a small number of copies (often just one copy) of a data block with a given content. Advanced deduplication can be enabled only if data reduction is enabled first.

The Advanced Deduplication switch is available only on:

- Dynamic or Traditional pools in Unity 380F, 480F, 680F, and 880F systems
- Dynamic pools in Unity All-Flash 450F, 550F, and 650F systems
- All-Flash pools and Hybrid Flash pools in Unity Hybrid 380, 480, 680, and 880 systems

Both All-Flash pools and Hybrid Flash pools can contain both data reduction-enabled and non-data reduction enabled LUNs, file systems, or datastores.

When data reduction and advanced deduplication are enabled, or were previously enabled and then disabled, for a storage resource in a hybrid pool, the percentage of Flash drives in the pool must be 10% or higher. Otherwise, an error will occur. For additional information, refer to the *Dell Unity: Data Reduction* white paper.

Managing the FAST Cache (supported physical deployments only)


Topics:

- [About the FAST Cache](#)
- [FAST Cache best practices](#)
- [Create the FAST Cache](#)
- [Enable or disable the FAST Cache for a pool](#)
- [View system-level FAST Cache information](#)
- [View pool-level Fast Cache Information](#)
- [Expand or shrink the FAST Cache capacity](#)
- [Delete the FAST Cache](#)

About the FAST Cache

FAST Cache functionality is supported for use with dynamic or traditional non-All-Flash pools on hybrid Flash arrays.

The FAST Cache is a large capacity secondary cache that uses SAS Flash drives to improve system performance by extending the existing caching capacity of the storage system. You can configure individual pools to use the FAST Cache. When you do this, all storage resources which are created in those pools use the FAST Cache.

 **NOTE:** The use of SAS Flash 3 or SAS Flash 4 type drives for the FAST Cache is not supported.

The FAST Cache provides the following benefits:

- **Improved Total Cost of Ownership (TCO)**—The FAST Cache reduces the number of drive spindles that are needed for a given level of performance and allows you to replace high-performance drives with more economical options. This is possible because the load on back-end hard drives is reduced. Because the FAST Cache copies the busy subsets of large storage capacities in 64 KB data chunks to Flash drives, the busiest areas of many LUNs or file systems use a small set of Flash drives. This allows a large set of slower, less expensive drives to deliver the performance typically provided by faster, more expensive drives. Over time, new pools with slower NL-SAS drives can replace pools that have faster SAS drives, while maintaining the same application performance. This improves the financial (IOPS/dollar) and power (IOPS/KWH) efficiency of the storage system.
- **Reduced complexity**—At a system level, the FAST Cache helps make the most efficient use of Flash drive capacity. The FAST Cache does this by using Flash drives for the most frequently accessed data in the storage system instead of dedicating Flash drives to a particular application. It adjusts to a hot spot anywhere in the storage system, so that you do not need to analyze specific application requirements. It provides better performance to all applications in the storage system while using fewer Flash drives.
- **Highly scalable cache capacity**—The FAST Cache can scale up to a larger capacity than the maximum DRAM Cache capacity. This larger cache capacity means that all or a higher percentage of a customer's working set can be promoted into the better performing FAST Cache. The working set is the most active part of an application.

FAST Cache drive configurations

The FAST Cache uses SAS Flash 2 drives that are configured in RAID 1 pairs. The drive sizes used to configure the FAST Cache and maximum FAST Cache capacity vary by model. For information, see the compatibility and interoperability documentation. Go to the support website to locate this documentation.

You cannot mix drive types or sizes in the FAST Cache.

FAST Cache actions

You can expand or shrink the FAST Cache by adding or removing drives. You can also delete the FAST Cache, but this can impact system performance and the performance of pools that are using the FAST Cache. The FAST Cache deletion process takes time, because the storage system has to flush all the cache data to the pool drives.

FAST Cache best practices

Use the FAST Cache under the following circumstances:

- Large working sets of data must be stored and read, accessed frequently, and readily available.
- The system has enough drives for both data storage and performance improvement.

For example, use the FAST Cache if the system has enough SAS Flash drives, is active with daily read/write transactions of big blocks of data, and is not used as long term storage for archival purposes that would be accessed on a weekly or monthly basis.

Avoid using the FAST Cache for data that is sequential or that uses large-block I/O. For example, do not use the FAST Cache for database logs or circular logs, since their data is sequential.

Create the FAST Cache

Prerequisites

To create the FAST Cache, the FAST Cache license must be installed on the system and the system must have available SAS Flash 2 drives in sizes supported by the model. (See the compatibility and interoperability documentation for this information).

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST Cache**.
2. Select **Create**.
3. Select the size and number of drives to use for the FAST Cache. The available drives are based on the number of available SAS Flash drives, spare drive policy, and system limits.
4. Optionally, clear **Enable FAST Cache for existing pools**, if you do not want the FAST Cache enabled for all existing pools. You can enable the FAST Cache for individual pools at a later time.

Enable or disable the FAST Cache for a pool

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and then select the **Edit** icon.
3. On the **General** tab, select or clear **Use FAST Cache**.

View system-level FAST Cache information

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST Cache** to view the status, total space, and number of drives in the FAST Cache.
2. Select the **Settings** icon, and then select **Storage Configuration > FAST Cache Disks** to view the name and size of each drive in the FAST Cache.

View pool-level Fast Cache Information

About this task

At the pool level, you can verify whether a specific pool is configured to use the FAST Cache.

Steps

1. Under **Storage**, select **Pools**.
2. Select the pool, and then select the **Edit** icon.
3. On the **General** tab, view whether the **Use FAST Cache** field is enabled.


Expand or shrink the FAST Cache capacity

About this task

If you have a small FAST Cache, and your business's growth requires larger blocks of data to be read and written, you can expand the FAST Cache capacity by adding SAS Flash 2 drives.

On rare occasions, you might want to shrink the FAST Cache capacity. For example, consider doing so under these circumstances:

- The FAST Cache is too large for users' applications.
- You need to repurpose some SAS Flash 2 drives to create storage pools with higher performance.
- The storage system does not have enough slots left to be populated with additional drives for storage.

 **NOTE:** Shrinking the FAST Cache can take a long time and can affect the performance of pools currently using the FAST Cache, because data on the drives being removed from the FAST Cache must be flushed to the pool drives.

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST Cache**.
2. To expand the FAST Cache, select **Expand**. Then select the drives to add to the FAST Cache.
3. To shrink the FAST Cache, select **Shrink**. Then select the drives to remove from the FAST Cache.

Delete the FAST Cache

About this task

Deleting the FAST Cache frees up drives for other uses. However, it can impact system performance and the performance of pools currently using the FAST Cache. The FAST Cache deletion process can take a long time, because it has to flush all of the cache data to the pool drives.

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST Cache**.
2. Select **Delete**.

Managing FAST VP (supported physical deployments only)

Topics:

- [About FAST VP](#)
- [Data relocation with FAST VP](#)
- [FAST VP tiering policy](#)
- [Enable or disable scheduled data relocations at the system level](#)
- [Stop a pool from participating in scheduled data relocations](#)
- [Manually initiate or stop data relocations for a pool](#)
- [Pause or resume data relocations for all pools](#)
- [Change system-level data relocation information](#)
- [Change data relocation information for a pool](#)
- [View pool-level FAST VP information](#)

About FAST VP

FAST VP functionality is supported for use with dynamic or traditional multi-tiered pools on hybrid Flash arrays.

Fully Automated Storage Tiering for Virtual Pools (FAST VP) enables the system to retain the most frequently accessed or important data on fast, high-performance drives, and move the less frequently accessed and less important data to lower-performance, cost-effective drives.

FAST VP does the following:

- Monitors the usage of the data in a tiered pool. Tiered pools are heterogeneous pools that are configured with multiple classes of drives (SAS Flash 2 and/or SAS Flash 3, plus SAS and/or NL-SAS).
- Depending on the tiering policy, uses the monitoring statistics to automatically relocate data chunks, at 256 MB granularity, to other tiers within the pool. For example, the Start High then Auto-Tier policy relocates data to the storage tier that is best suited for that data, based on relative activity.
- Performs load balancing across the drives in tiered and nontiered pools.

FAST VP is an automated feature that optimizes drive utilization. It requires little manual intervention.

NOTE:

- If you want to configure and use the FAST VP feature, the FAST VP license must be installed on the system. All-Flash models of Unity do not support FAST VP.
- FAST VP can use all supported drive types except for SAS Flash 4.

The data relocation that FAST VP performs can help you achieve the following benefits:

- **Increased performance**—Sometimes, you can double performance throughput by adding less than 10 percent of a pool's total capacity in SAS Flash 2 or SAS Flash 3 drives.
- **Reduced Total Cost of Ownership (TCO)**—Using a combination of NL-SAS, SAS, SAS Flash 2, and SAS Flash 3 drives instead of all SAS drives enables you to address performance requirements and still reduce the drive count. Sometimes, you can achieve up to a two-thirds reduction in drive count by using FAST VP.

Data relocation with FAST VP

Data relocation (also called auto-tiering) is the process by which the system relocates data to another tier within the same pool and rebalances data within a tier to achieve maximum pool performance gain. Data relocation can be scheduled or manually initiated.

The data relocation process does the following:


- Relocates the most active data up to the highest available tier.
- Relocates the least active data to lower tiers, as needed to free up space for more active data in the higher tiers.
- Rebalances data within a tier.

Scheduled data relocation

Scheduled data relocation is enabled for all pools by default when FAST VP is installed on your system. The system performs scheduled data relocations for pools participating in the automatic data relocation schedule. By default, the data relocation window is every day from 22:00 to 6:00 (UTC time) the next day. The specified data relocation times provide the window of time in which data relocations occur on an ongoing basis, if the system determines the relocations are needed.

You can do the following to configure scheduled data relocations:

- Disable or re-enable scheduled data relocations for the system.
- Change the days on which data relocation occurs and the data relocation window start time and end time for each scheduled day.
- Configure individual pools to participate in scheduled data relocations.
- Pause all scheduled data relocations. (This action also pauses manual data relocations that are currently in progress.)
- Resume all scheduled data relocations (This action also resumes paused manual data relocations whose specified end times have not expired.)

 **NOTE:** The system rebalances data within a tier automatically when a pool tier is expanded.

Manual data relocation

You can manually initiate data relocations for individual pools. When you start a manual data relocation, you specify a data relocation rate and an end time for the data relocation operation.

You can do the following to control manual data relocations:

- Start or stop a manual data relocation.
- Pause manual relocations that are currently in progress. This action also pauses all scheduled data relocations.
- Resume paused manual relocations whose end times are not over. This action also resumes all scheduled data relocations.

FAST VP tiering policy

The following table describes the FAST VP tiering policy settings, which are defined at the data-resource level. This policy defines both the initial tier placement and the ongoing automated tiering of data during data relocation operations.

Table 4. FAST VP tiering policy settings

Tiering policy	Corresponding initial tier placement	Description
Start High then Auto-Tier (default)	Highest available tier	Recommended setting. Sets the initial data placement to the highest-performing drives with available space, and then relocates portions of the storage resource's data based on I/O activity.
Auto-Tier	Optimize for pool performance	Sets the initial data placement to an optimum, system-determined setting, and then relocates portions of the storage resource's data based on the storage resource's performance statistics such that data is relocated among tiers according to I/O activity.
Highest Available Tier	Highest available tier	Sets the initial data placement and subsequent data relocation (if

Table 4. FAST VP tiering policy settings (continued)


Tiering policy	Corresponding initial tier placement	Description
		applicable) to the highest-performing drives with available space.
Lowest Available Tier	Lowest available tier	Sets the initial data placement and subsequent data relocation (if applicable) to the most cost-effective drives with available space.

Enable or disable scheduled data relocations at the system level

About this task

When you disable scheduled data relocations at the system level, no pools in the system will participate in scheduled data relocations. If you disable scheduled data relocations at the system level when a data relocation is taking place, the system will complete the data relocation.

When you enable scheduled data relocations at the system level, pools that have scheduled data relocations enabled at the pool level will participate in scheduled data relocations.

 **NOTE:** Scheduled data relocations for the storage system are enabled by default.

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST VP**.
2. Do either of the following:
 - Select **Pause** to disable data relocations for the storage system.
 - Select **Resume** to enable data relocations for the storage system, and optionally select **Modify data relocation schedule** to change the data relocation schedule.

Stop a pool from participating in scheduled data relocations

About this task

If you stop a pool from participating in scheduled data relocations when a data relocation is taking place, the system will complete the data relocation.

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and click the **Edit** icon.
3. Select **Fast VP**.
4. If the **Data relocations scheduled** field has a value of No, select the **Edit** icon, and select **Yes** on the confirmation message.

Manually initiate or stop data relocations for a pool

About this task

You can manually initiate or stop data relocations for pools.

Steps

1. Under **Storage**, select **Pools**.
2. Select the pool for which you want to initiate or stop data relocations, and then select the **Edit** icon.
3. Select **FAST VP**.
4. Do either of the following:
 - To start data relocations for the pool, choose **Start Relocation**, and select the rate and end time for the data relocation. A data relocation rate of **Low** has the least impact on system performance, **Medium** (the default) has a moderate impact on system performance, and **High** has the most impact on system performance.
 - To stop data relocations, choose **Stop Relocation**. This stops all data relocations for the pool, including manually initiated data relocations that are currently in progress and scheduled data relocations.

Pause or resume data relocations for all pools

About this task

Pausing data relocations stops all scheduled data relocations as well as manual data relocations that are in progress. Resuming data relocations restarts these data relocations.

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST VP**.
2. Pause or resume data relocations.

Change system-level data relocation information

About this task

System-level data relocation information includes whether scheduled relocations are enabled for the system. If scheduled data relocations are enabled, the information includes:

- Total amount of data in the system scheduled to move to higher and lower tiers.
- Total amount of data in the system to be rebalanced within a tier
- Estimated time needed for the scheduled data relocations.

Steps

1. Select the **Settings** icon, and then select **Storage Configuration > FAST VP**.
2. Optionally, take one of the following actions:
 - Select **Pause** (if enabled) to pause all scheduled data relocations and manual data relocations that are currently in progress.
 - Select **Resume** (if enabled) to resume paused data relocations.
 - Modify the data relocation rate. A data relocation rate of Low has the least impact on system performance, Medium (the default) has a moderate impact on system performance, and High has the most impact on system performance.
 - Select **Modify data relocation schedule** to change the data relocation schedule.

Change data relocation information for a pool

About this task

You can specify whether a pool is scheduled to participate in scheduled data relocations. You can also manually start or stop data relocations on the pool.

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and click **Edit**.
3. Select **FAST VP**.

4. Make the appropriate changes.

View pool-level FAST VP information

About this task

If a FAST VP license is installed on the system, you can view the following information for a pool:

- Whether the pool participates in scheduled data relocations.
- Estimated time needed for scheduled data relocations.
- Start and end time for the most recent data relocation.
- Number and type of drives in each tier.
- Amount of data in the pool scheduled to move to higher and lower tiers.
- Amount of data in the pool scheduled to be rebalanced within a tier.

Steps

1. Under **Storage**, select **Pools**.
2. Select the relevant pool, and then select the **Edit** icon.
3. Select the **FAST VP** tab.
4. Optionally click **Manage FAST VP system settings** to view the data relocation rate.
 - **Low** has the least impact on system performance.
 - **Medium** (the default) has a moderate impact on system performance.
 - **High** has the most impact on system performance.

Understanding RAID (physical deployments only)

Topics:

- [About RAID](#)
- [RAID configurations](#)
- [RAID for dynamic pools in Unisphere](#)
- [Drive IOPS by RAID type](#)

About RAID

Redundant Array of Independent Disks (RAID) is a method for providing high levels of storage reliability by arranging drives in groups, and dividing and replicating data among the drives in a group. You select drive types and RAID configurations (RAID types and stripe widths) when you create a pool. Once a pool is configured, you cannot change the RAID type of a tier. However, in hybrid models, you can add a new tier with a different RAID type. (All-Flash models have only one tier.)

RAID usually has the characteristics of parity, striping, or both:

- Parity provides redundancy for blocks of data on the drives. Depending on the RAID type, this provides the ability to continue to operate with the loss of one or more drives.
- Striping provides a mechanism for processing data that allows the comprehensive read/write performance of a RAID group to exceed the performance of its component drives.

In hybrid models that are licensed for FAST VP and that have unused drives of different types, you can configure multiple storage tiers for the pool and can specify a different RAID configuration for each tier.

The system supports RAID 5, 6, and 1/0 (also called RAID 10).

RAID in dynamic pools

In a dynamic pool, drives are partitioned into drive extents. These extents are combined into RAID extents, which are spread across multiple drives. The number of drive extents within a RAID extent depends on the RAID type and width. For example, a RAID 5 (4+1) RAID extent contains five drive extents. For redundancy purposes, the system ensures that no RAID extent contains two drive extents from the same drive. Also, each drive extent can only be part of one RAID extent.

A RAID group in a dynamic pool is a collection of RAID extents. The number and size of RAID groups in a dynamic pool can vary depending on the number of drives and how the pool was created and expanded. A RAID extent can only be part of one RAID group.

The following illustration shows a RAID group with a RAID 5 (4 + 1) configuration. This configuration requires six drives to meet the requirement for spare space. In this illustration, each rectangle represents one drive extent, and each color represents one RAID extent:

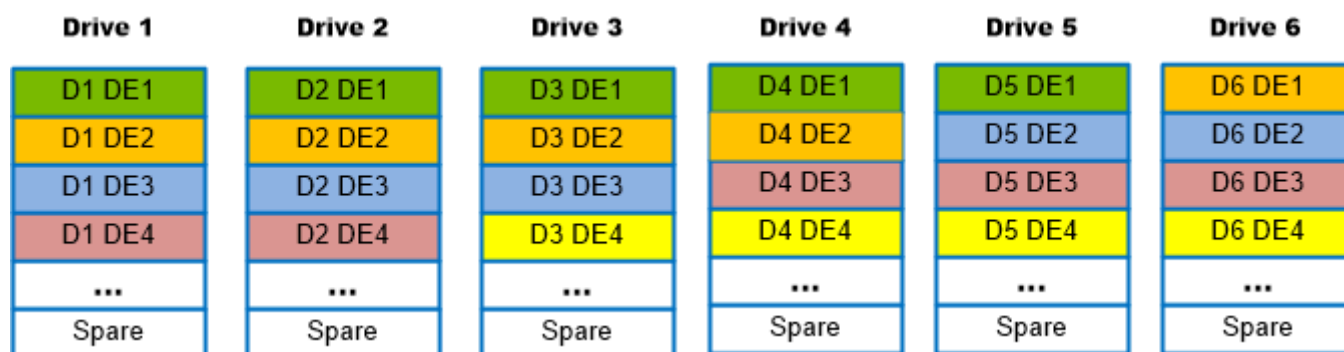


Figure 1. RAID 5 (4+1) group in a dynamic pool

RAID in traditional pools

In a traditional pool, a RAID group contains sets of drives with the same capacity and redundancy on which you create one or more storage resources. The system creates one or more RAID groups for the pool based on the specified configuration. For example, when you create a storage resource in a RAID 5 (4 + 1) group, data is distributed equally across the five drives in the RAID group. If you want to create a pool using more than five drives, you must do so in multiples of the selected RAID stripe width.

The following illustration shows a RAID 5 (4 + 1) group, along with the required spare drive.

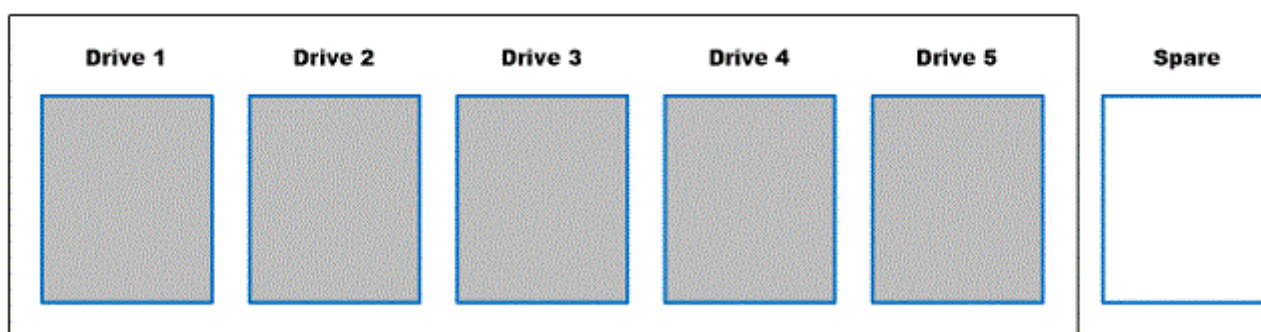


Figure 2. RAID 5 (4 + 1) group plus one spare drive in a traditional pool

RAID configurations

Pool tiers are built using a set of one or more individual drive groups based on the tier's RAID type and stripe width. The RAID type determines the performance characteristics of each drive group. The stripe width determines the fault characteristics of each drive group.

For example, a RAID 5 drive group can still operate with the loss of one drive (traditional pools) or it's equivalent (dynamic pools). A RAID 5 (4+1), 5 drive configuration has less risk of multiple drive faults than a RAID 5 (12+1), 13 drive configuration.

The following table describes the supported RAID types for the intended storage usage:

Table 5. Supported RAID levels

RAID level	Description
RAID 1/0 (also called RAID 10)	Best suited for applications with fast or high processing requirements, such as enterprise servers and moderate-sized database systems. Provides both high performance and reliability at medium cost, while providing lower capacity per drive.
RAID 5	Best suited for transaction processing and often used for general purpose storage, as well as for relational database and enterprise resource systems. Depending on the drives used, this RAID type can provide a fairly low cost per MB while still retaining redundancy.

Table 5. Supported RAID levels (continued)

RAID level	Description
RAID 6	Best suited for read-biased workloads, such as archiving and backup to drive.

Mixed RAID configurations (traditional pools only)

If FAST VP is installed on the system, you can create a pool with multiple storage tiers. Each tier can have its own RAID type. Only one RAID type can be used within a tier, but the tier can have different stripe configurations. For example, you can mix RAID 5 (4+1) and RAID 5 (8+1) in a tier.

To do this:

- Select the **Maximum Capacity** RAID configuration when you create the pool. This configuration might mix RAID types in the pool.
- Expand the pool using a different stripe width than currently exists in the pool.

RAID for dynamic pools in Unisphere

Unisphere automatically assigns a RAID width to a dynamic pool based on the selected RAID type and the selected number of drives. The following table describes how the RAID width gets assigned:

RAID level	Number of drives	Number of drives (if spare number is 2)	RAID Width
RAID 5	6 to 7	7 to 8	4+1
	10 to 13	11 to 14	8+1
	14 or more	15 or more	12+1
RAID 6	7 to 8	8 to 9	4+2
	9 to 10	10 to 11	6+2
	11 to 12	12 to 13	8+2
	13 to 14	14 to 15	10+2
	15 to 16	16 to 17	12+2
	17 or more	18 or more	14+2
RAID 1/0	3 to 4	4 to 5	1+1
	5 to 6	6 to 7	2+2
	7 to 8	8 to 9	3+3
	9 or more	10 or more	4+4

Drive IOPS by RAID type

Front-end application workloads translate into different back-end drive workloads based on the RAID type in use. For front-end reads, there is no impact by RAID type: 1 front-end read I/O equals 1 back-end read I/O.

The following table shows the impact by RAID type for random front-end writes.

Table 6. IOPS by RAID type for front-end writes

RAID type	IOPS per 1 front-end write I/O
RAID 1/0	2 back-end write I/Os
RAID 5	2 back-end reads and 2 back-end writes

Table 6. IOPS by RAID type for front-end writes (continued)

RAID type	IOPS per 1 front-end write I/O
RAID 6	3 back-end reads and 3 back-end writes

Configuring pools using the CLI

Topics:

- [Configure pools automatically](#)
- [Configure custom pools](#)
- [Manage pool tiers](#)
- [View pool resources](#)
- [View storage profiles \(physical deployments only\)](#)
- [Manage drive groups \(physical deployments only\)](#)
- [Manage storage system capacity settings](#)
- [Manage system tier capacity settings](#)
- [Change disk settings \(virtual deployments only\)](#)
- [Pool use cases](#)
- [The show action command](#)

Configure pools automatically

Pools are the groups of physical drives or virtual disks, called disk groups, on which you create storage resources. The system can automatically configure pools by selecting the appropriate disk groups based on the type and availability of drives in the system.

In order to automatically create pools, the following criteria must be met:

- A FAST VP/FAST Cache license is not installed
- No other pools have been created
- For virtual deployments, all virtual disks have been assigned a tier type

i **NOTE:** [Configure custom pools](#) explains how to configure custom pools. Before you create storage resources (including NAS servers), you must configure at least one pool.

i **NOTE:** All disks on the system will be used for storage pools.

The following table lists the attributes for automatic pool configuration.

Table 7. Automatic pool configuration attributes


Attribute	Description
Target	Type of drive configuration. Value is one of the following: <ul style="list-style-type: none"> • <code>pool</code> - Drives configured in a pool. • <code>spares</code> - Drives assigned to pools as spares. The number of spares assigned to a pool is dependent on the drive type and pool type: <ul style="list-style-type: none"> ○ For Capacity pools, no spare drives are assigned. ○ For Performance pools, a spare drive is assigned for the first 0-30 drives, and then another spare drive is assigned for every group of thirty drives after that. ○ For Flash pools, a spare drive is assigned for the first 0-30 drives, and then another spare drive is assigned for every group of thirty drives after that.
Name	Name of the pool. The system allocates drives to one or more of the following pools based on the types and characteristics of the disks on the system: <ul style="list-style-type: none"> • <code>Capacity</code> - Storage allocated from near-line (NL) serial attached SCSI (SAS) drives. Provides high-capacity storage, but with lower overall performance to regular SAS and Enterprise Flash Drive (EFD) drives. Use NL SAS drives to provide extremely economical storage for operations, such as data backup, that do not require high I/O performance.

Table 7. Automatic pool configuration attributes (continued)

Attribute	Description
	<ul style="list-style-type: none">• Performance - Storage allocated from serial attached SCSI (SAS) disks. Provides medium performance and medium capacity storage for applications that require balance of performance and capacity.• Flash - Storage allocated from Flash drives. Extremely high level performance, but at a relatively high cost per GB of storage. Flash drives are most applicable to applications that require high I/O performance and energy efficiency. Depending on the pool type, the system configures the drives into different RAID groups and assigns drives to pools as spares. The Unisphere online help provides more details about pools and spares.
Drives (current)	List of drives currently in the pool.
Drives (new)	List of drives to be added to the pool.
RAID level (physical deployments only)	RAID level applied.
Stripe length (physical deployments only)	Comma-separated list of drives in the stripe.

Initiate automatic pool configuration

Start configuring pools automatically. [View configuration settings for automatic pool creation](#) displays the configuration settings that the system will apply when you run this command.

 **NOTE:** All disks on the system will be used to configure storage pools automatically. You must explicitly confirm this operation to proceed.

Format

```
/stor/config/auto set
```

Action qualifier

Qualifier	Description
-async	Run action in asynchronous mode.

Example

The following command initiates automatic pool configuration:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/auto set
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
All disks in the system will be used to configure pools. Do you want to proceed?
yes / no: yes
Operation completed successfully.
```

View configuration settings for automatic pool creation

View the settings for automatic pool creation that will be applied to the system. [Initiate automatic storage pool configuration](#) explains how to apply these settings to the system.

 **NOTE:** The `show action` command explains how to change the output format.

Format

```
/stor/config/auto show
```

Example (physical deployments only)

The following command shows how pools and spares will be configured automatically on the system:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/auto show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      Target      = Pool
      Name        = Performance
      Drives (current) = 5 x 600GB SAS; 5 x 300GB SAS
      Drives (new)    = 5 x 600GB SAS
      RAID level     = 5
      Stripe length  = 5,9

2:      Target      = Pool
      Name        = Capacity
      Drives (current) = 10 x 1TB NL-SAS
      Drives (new)    = 2 x 1TB NL SAS
      RAID level     = 5
      Stripe length  = 5,9


3:      Target      = Pool
      Name        = Extreme Performance
      Drives (current) = 10 x 100GB EFD
      Drives (new)    = 10 x 100GB EFD
      RAID level     = 1
      Stripe length  = 2

4:      Target      = Spares
      Name        = Unused / Hot Spare Candidates
      Drives (current) = 1 x 600GB SAS; 1 x 300GB SAS; 1 x 1TB NL SAS
      Drives (new)    = 1 x 100GB EFD
      RAID level     =
      Stripe length  =
```

Configure custom pools

Pools are the groups of drives on which you create storage resources. Configure pools based on the type of storage resource and usage that will be associated with the pool, such as file system storage optimized for database usage. The storage characteristics differ according to the following:

- Type of drive used to provide the storage.
- (dual-SP virtual deployments only) RAID level implemented for the storage.

 **NOTE:** Before you create storage resources, you must configure at least one pool.

The following table lists the attributes for pools:

Table 8. Custom pool attributes

Attribute	Description
ID	ID of the pool.
Name	Name of the pool.

Table 8. Custom pool attributes (continued)

Attribute	Description
Type	Pool type. Valid values are: <ul style="list-style-type: none"> • Dynamic • Traditional
Description	Brief description of the pool.
Total space	Total storage capacity of the pool.
Current allocation	Amount of storage in the pool allocated to storage resources.
Preallocated space	Amount of storage space reserved in the pool by storage resources for future needs to make writes more efficient. The pool may be able to reclaim some of this space if total pool space is running low. This value equals the sum of the <code>sizePreallocated</code> values of each storage resource in the pool.
Remaining space	Amount of storage in the pool not allocated to storage resources.
Subscription	For thin provisioning, the total storage space subscribed to the pool. All pools support both standard and thin provisioned storage resources. For standard storage resources, the entire requested size is allocated from the pool when the resource is created, for thin provisioned storage resources only incremental portions of the size are allocated based on usage. Because thin provisioned storage resources can subscribe to more storage than is actually allocated to them, pools can be over provisioned to support more storage capacity than they actually possess. <i>i</i> NOTE: The system automatically generates an alert when the total pool usage reaches 85% of the pool's physical capacity. <code>-alertThreshold</code> specifies the alert threshold value.
Subscription percent	For thin provisioning, the percentage of the total space in the pool that is subscription storage space.
Alert threshold	Threshold for the system to send an alert when hosts have consumed a specific percentage of the subscription space. Value range is 50 to 85.
Drives	List of the types of drives on the system, including the number of drives of each type, in the pool. If FAST VP is installed, you can mix different types of drives to make a tiered pool. However, SAS Flash 4 drives must be used in a homogeneous pool.
Number of drives	Total number of drives in the pool.
Number of unused drives	Number of drives in the pool that are not being used.
RAID level (physical deployments only)	RAID level of the drives in the pool.
Stripe length (physical deployments only)	Number of drives the data is striped across.
Rebalancing	Indicates whether a pool rebalancing is in progress. Valid values are: <ul style="list-style-type: none"> • yes • no
Rebalancing progress	Indicates the progress of the pool rebalancing as a percentage.
System defined pool	Indication of whether the system configured the pool automatically. Valid values are: <ul style="list-style-type: none"> • yes • no
Health state	Health state of the pool. The health state code appears in parentheses. Valid values are: <ul style="list-style-type: none"> • Unknown (0) —Health is unknown. • OK (5) —Operating normally.

Table 8. Custom pool attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> • <code>OK BUT (7)</code> —Pool has exceeded its user-specified threshold or the system specified threshold of 85%. • <code>Degraded/Warning (10)</code> —Pool is operating, but degraded due to one or more of the following: <ul style="list-style-type: none"> ◦ Pool has exceeded the user-specified threshold. ◦ Pool is nearing capacity. ◦ Pool is almost full. ◦ Pool performance has degraded. • <code>Major failure (20)</code> —Dirty cache has made the pool unavailable. • <code>Critical failure (25)</code> —Pool is full. To avoid data loss, add more storage to the pool, or create more pools. • <code>Non-recoverable error (30)</code> —Two or more drives in the pool have failed, possibly resulting in data loss.
<code>Health details</code>	Additional health information. See Appendix A, Reference, for health information details.
<code>FAST Cache enabled (physical deployments only)</code>	Indicates whether FAST Cache is enabled on the pool. Valid values are: <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code>
<code>Non-base size used</code>	Quantity of storage used for thin clone and snapshot data.
<code>Auto-delete state</code>	Indicates the state of an auto-delete operation on the pool. Valid values are: <ul style="list-style-type: none"> • <code>Idle</code> • <code>Running</code> • <code>Could not reach LWM</code> • <code>Could not reach HWM</code> <p>NOTE: If the auto-delete operation cannot satisfy the high water mark, and there are snapshots in the pool, the auto-delete operation sets the auto-delete state for that watermark to <code>Could not reach HWM</code>, and generates an alert.</p> <ul style="list-style-type: none"> • <code>Failed</code>
<code>Auto-delete paused</code>	Indicates whether an auto-delete operation is paused. Valid values are: <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code>
<code>Auto-delete pool full threshold enabled</code>	Indicates whether the system will check the pool full high water mark for auto-delete. Valid values are: <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code>
<code>Auto-delete pool full high water mark</code>	The pool full high watermark on the pool.
<code>Auto-delete pool full low water mark</code>	The pool full low watermark on the pool.
<code>Auto-delete snapshot space used threshold enabled</code>	Indicates whether the system will check the snapshot space used high water mark for auto-delete. Valid values are: <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code>
<code>Auto-delete snapshot space used high water mark</code>	High watermark for snapshot space used on the pool.
<code>Auto-delete snapshot space used low water mark</code>	Low watermark for snapshot space used on the pool.

Table 8. Custom pool attributes (continued)

Attribute	Description
Data Reduction space saved (physical deployments only)	Storage size saved on the pool by using data reduction. i NOTE: Data reduction is available for thin LUNs and thin file systems. The thin file systems must be created on Unity systems running version 4.2.x or later.
Data Reduction percent (physical deployments only)	Storage percentage saved on the pool by using data reduction. i NOTE: Data reduction is available for thin LUNs and thin file systems. The thin file systems must be created on Unity systems running version 4.2.x or later.
Data Reduction ratio (physical deployments only)	Ratio between data without data reduction and data after data reduction savings. i NOTE: Data reduction is available for thin LUNs and thin file systems. The thin file systems must be created on Unity systems running version 4.2.x or later.
All flash pool	Indicates whether the pool contains only Flash drives. Valid values are: <ul style="list-style-type: none"> • yes • no

Create pools

Create a dynamic or traditional pool.

- When creating a dynamic pool with NL-SAS drives, RAID 5, 6, and 10 are allowed.
- Both traditional pools and dynamic pools are supported in the CLI and REST API for Unity All-Flash models running OE version 4.2.x or later, and for Unity hybrid models running OE version 5.2.x or later. The default pool type is dynamic.
- Traditional pools are supported in all Unity hybrid and virtual models. They are also supported in Unity All-Flash models running OE version 4.1.x or earlier.
- The minimum number of drives to use for a dynamic pool is the selected RAID width plus one drive.

Format

```
/stor/config/pool create [-async] -name <value> [-type {dynamic [-hotSpareCapacity <value>] | traditional}] [-descr <value>] [-diskGroup <value> -drivesNumber <value> [-storProfile <value>] | -disk <value>] [-tier <value>] [-alertThreshold <value>] [-snapPoolFullThresholdEnabled {yes|no}] [-snapPoolFullHWM <value>] [-snapPoolFullLWM <value>] [-snapSpaceUsedThresholdEnabled {yes|no}] [-snapSpaceUsedHWM <value>] [-snapSpaceUsedLWM <value>]
```

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode. i NOTE: Simultaneous commands, asynchronous or synchronous, may fail if they conflict in trying to manage the same system elements.
-name	Type a name for the pool.
-type	(Available only for systems that support dynamic pools) Specify the type of pool to create. Value is one of the following: <ul style="list-style-type: none"> • dynamic <ul style="list-style-type: none"> ◦ -hotSpareCapacity Specifies the hot spare capacity for every 32 drives of dynamic pools. Valid values are 1 (reserve 1 drive capacity for


Qualifier	Description
	<p>every 32 drives as a spare) or 2 (reserve 2 drives capacity for every 32 drives as a spare).</p> <p>i NOTE: The <code>-hotSpareCapacity</code> setting is used for a tier. This means that all disk groups in the same tier share the same <code>-hotSpareCapacity</code> value.</p> <ul style="list-style-type: none"> • <code>traditional</code> <p>Default value is <code>dynamic</code>.</p>
<code>-descr</code>	Type a brief description of the pool.
<code>-storProfile</code> (physical deployments only)	Type the ID of the storage profiles, separated by commas, to apply to the pool, based on the type of storage resource that will use the pool and the intended usage of the pool. View storage profiles (physical deployments only) explains how to view the IDs of available storage profiles on the system. If this option is not specified, a default RAID configuration is selected for each particular drive type in the selected drive group: NL-SAS (RAID 6 with a stripe length of 8), SAS (RAID 5 with a stripe length of 5), or Flash (RAID 5 with a stripe length of 5).
<code>-diskGroup</code> (physical deployments only)	Type a comma-separated list of IDs of the drive groups to use in the pool. Specifying drive groups with different drive types causes the creation of a multi-tier pool. View drive groups explains how to view the IDs of the drive groups on the system.
<code>-drivesNumber</code> (physical deployments only)	Specify the drive numbers, separated by commas, from the selected drive groups to use in the pool. If this option is specified when <code>-storProfile</code> is not specified, the operation may fail when the <code>-drivesNumber</code> value does not match the default RAID configuration for each drive type in the selected drive group.
	<p>i NOTE: When creating a dynamic pool on a hybrid Flash array, you cannot specify HDD system drives. When creating a traditional pool, you can specify HDD system drives.</p>
<code>-disk</code> (virtual deployments only)	Specify the list of drive IDs, separated by commas, to use in the pool. Specified drives must be reliable storage objects that do not require additional protection.
<code>-tier</code> (virtual deployments only)	Specify the comma-separated list of tiers to which the drives are assigned. If the tier is omitted, it is assigned automatically if tiering information for the associated drive is available. Value is one of the following: <ul style="list-style-type: none"> • <code>capacity</code> • <code>performance</code> • <code>extreme</code>
<code>-alertThreshold</code>	For thin provisioning, specify the threshold, as a percentage, when the system will alert on the amount of subscription space used. When hosts consume the specified percentage of subscription space, the system sends an alert. Value range is 50% to 85%.
<code>-FASTCacheEnabled</code> (physical deployments only)	Specify whether to enable FAST Cache on the pool. Value is one of the following: <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code> Default value is <code>yes</code> .
<code>-snapPoolFullThresholdEnabled</code>	Indicate whether the system should check the pool full high water mark for auto-delete. Value is one of the following: <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code> Default value is <code>yes</code> .

Qualifier	Description
-snapPoolFullHWM	Specify the pool full high watermark for the pool. Valid values are 1-99. Default value is 95.
-snapPoolFullLWM	Specify the pool full low watermark for the pool. Valid values are 0-98. Default value is 85.
-snapSpaceUsedThresholdEnabled	Indicate whether the system should check the snapshot space used high water mark for auto-delete. Value is one of the following: <ul style="list-style-type: none"> • yes • no Default value is yes.
-snapSpaceUsedHWM	Specify the snapshot space used high watermark to trigger auto-delete on the pool. Valid values are 1-99. Default value is 95.
-snapSpaceUsedLWM	Specify the snapshot space used low watermark to trigger auto-delete on the pool. Valid values are 0-98. Default value is 20.

 **NOTE:** Use the [Change disk settings \(virtual deployments only\)](#) command to change the assigned tiers for specific drives.

Example 1 (physical deployments only)

The following command creates a dynamic pool. This example uses storage profiles profile_1 and profile_2, six drives from drive group dg_2, and ten drives from drive group dg_28. The configured pool receives ID pool_2.

 **NOTE:** Before using the `stor/config/pool create` command, use the `/stor/config/profile show` command to display the dynamic pool profiles and the `/stor/config/dg show` command to display the drive groups.


```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! uemcli /stor/config/pool create
-name MyPool -descr "dynamic pool" -diskGroup dg_2,dg_28 -drivesNumber 6,10 -storProfile
profile_1,profile_2
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = pool_2
Operation completed successfully.
```

Example 2 (physical deployments only)

The following command creates a traditional pool in models that support dynamic pools. This example uses storage profiles tprofile_1 and tprofile_2, five drives from drive group dg_3, and nine drives from drive group dg_28. The configured pool receives ID pool_6.

 **NOTE:** Before using the `stor/config/pool create` command, use the `/stor/config/profile -traditional show` command to display the traditional pool profiles (which start with "t") and the `/stor/config/dg show` command to display the drive groups.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name
MyPool -descr "traditional pool" -diskGroup dg_3,dg_28 -drivesNumber 5,9 -storProfile
tprofile_1,tprofile_2 -type traditional
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = pool_6
Operation completed successfully.
```


Example 3 (physical deployments only)

The following command creates a traditional pool in models that do not support dynamic pools. This example uses storage profiles `profile_19` and `profile_20`, five drives from drive group `dg_15`, and nine drives from drive group `dg_16`. The configured pool receives ID `pool_5`.

NOTE: Before using the `stor/config/pool create` command, use the `/stor/config/profile show` command to display the traditional pool profiles and the `/stor/config/dg show` command to display the drive groups.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name
MyPool -descr "my big pool" -storProfile profile_19,profile_20 -diskGroup dg_15,dg_16
-drivesNumber 5,9 -FASTCacheEnabled yes
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = pool_5
Operation completed successfully.
```

Example 4 (virtual deployments only)

The following command creates a traditional pool with two virtual drives, `vdisk_0` and `vdisk_2` in the Extreme Performance tier. The configured pool receives ID `pool_4`.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name vPool
-descr "my virtual pool" -disk vdisk_0,vdisk_2
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = pool_4
Operation completed successfully.
```

Example 5 (physical deployments only)

The following command creates a dynamic pool, using 7 drives from disk group `dg_21` with a hot spare capacity of 1 drive; 6 drives from disk group `dg_31` with a hot spare capacity of 2 drives; and 7 drives from drive group `dg_18` with a hot spare capacity of 2 drives:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name pool
-diskGroup dg_21,dg_31,dg_18 -drivesNumber 7,6,7 -type dynamic -hotSpareCapacity 1,2,2
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
ID = pool_1
Operation completed successfully.
```

NOTE: The `-hotSpareCapacity` setting is used for a tier. This means that all disk groups in the same tier share the same `-hotSpareCapacity` value. In this example, drive groups `dg_31` and `dg_18` belong to the same tier, and share the same `-hotSpareCapacity` setting of 2.

Change pool settings

Change the subscription alert threshold, FAST Cache, and snapshot threshold settings for a pool.


Format

```
/stor/config/pool {-id <value> | -name <value>} set [-async] -name <value> [-  
descr <value>] [-alertThreshold <value>] [-snapPoolFullThresholdEnabled {yes|no}] [-  
snapPoolFullHWM <value>] [-snapPoolFullLWM <value>] [-snapSpaceUsedThresholdEnabled {yes|  
no}] [-snapSpaceUsedHWM <value>] [-snapSpaceUsedLWM <value>] [-snapAutoDeletePaused no]
```

Object qualifiers

Qualifier	Description
-id	Type the ID of the pool to change.
-name	Type the name of the pool to change.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.  NOTE: Simultaneous commands, asynchronous or synchronous, may fail if they conflict in trying to manage the same system elements.
-name	Type a name for the pool.
-descr	Type a brief description of the pool.
-alertThreshold	For thin provisioning, specify the threshold, as a percentage, when the system will alert on the amount of subscription space used. When hosts consume the specified percentage of subscription space, the system sends an alert. Value range is 50% to 84%.
-FASTCacheEnabled (physical deployments only)	Specify whether to enable FAST Cache on the pool. Value is one of the following: <ul style="list-style-type: none">• yes• no
-snapPoolFullThresholdEnabled	Indicate whether the system should check the pool full high water mark for auto-delete. Value is one of the following: <ul style="list-style-type: none">• yes• no
-snapPoolFullHWM	Specify the pool full high watermark for the pool. Valid values are 1-99. Default value is 95.
-snapPoolFullLWM	Specify the pool full low watermark for the pool. Valid values are 0-98. Default value is 85.
-snapSpaceUsedThresholdEnabled	Indicate whether the system should check the snapshot space used high water mark for auto-delete. Value is one of the following: <ul style="list-style-type: none">• yes• no
-snapSpaceUsedHWM	Specify the snapshot space used high watermark to trigger auto-delete on the pool. Valid values are 1-99. Default value is 95.
-snapSpaceUsedLWM	Specify the snapshot space used low watermark to trigger auto-delete on the pool. Valid values are 0-98. Default value is 20.
-snapAutoDeletePaused	Specify whether to pause snapshot auto-delete. Typing no resumes the auto-delete operation.

Example

The following command sets the subscription alert threshold for pool pool_1 to 70%:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool -id pool_1 -set -alertThreshold 70 -FASTCacheEnabled no
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = pool_1
Operation completed successfully.
```

Add drives to pools

Add new drives to a pool to increase its storage capacity.

NOTE:

- The minimum number of drives to use for a dynamic pool is the selected RAID width plus one drive.
- To extend a hybrid pool when data reduction and advanced deduplication are enabled for a storage resource in the pool, the percentage of Flash drives in the pool must be 10% or higher.


Format

```
/stor/config/pool {-id <value> | -name <value>} extend [-async] {-diskGroup <value> -drivesNumber <value> [-storProfile <value>] | -disk <value> [-tier <value>]}
```

Object qualifiers

Qualifier	Description
-id	Type the ID of the pool to extend.
-name	Type the name of the pool to extend.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.
-diskGroup (physical deployments only)	Type the IDs of the drive groups, separated by commas, to add to the pool.
-drivesNumber (physical deployments only)	Type the number of drives from the specified drive groups, separated by commas, to add to the pool. If this option is specified when -storProfile is not specified, the operation may fail when the -drivesNumber value does not match the default RAID configuration for each drive type in the selected drive group.  NOTE: When creating a dynamic pool on a hybrid Flash array, you cannot specify HDD system drives. When creating a traditional pool, you can specify HDD system drives.
-storProfile (physical deployments only)	Type the IDs of the storage profiles, separated by commas, to apply to the pool. If this option is not specified, a default RAID configuration is selected for each particular drive type in the selected drive group: <ul style="list-style-type: none">• NL-SAS (RAID 6 with a stripe length of 8)• SAS (RAID 5 with a stripe length of 5)• Flash (RAID 5 with a stripe length of 5)

Qualifier	Description
-disk (virtual deployments only)	Specify the list of drives, separated by commas, to add to the pool. Specified drives must be reliable storage objects that do not require additional protection.
-tier (virtual deployments only)	Specify the comma-separate list of tiers to which the added drives will be assigned. If a tier is not specified in this list, it will be assigned automatically if tiering information for the associated drive is available. Valid values are: <ul style="list-style-type: none"> • capacity • performance • extreme

Example 1 (physical deployments only)

The following command extends pool pool_1 with seven drives from drive group DG_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool -id pool_1 extend -diskGroup dg_1 -drivesNumber 7 -storProfile profile_12
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = pool_1
Operation completed successfully.
```

Example 2 (virtual deployments only)

The following command extends pool pool_1 by adding two virtual disks, vdisk_1 and vdisk_5.

```
uemcli -d 10.0.0.2 -u Local/joe -p MyPassword456! /stor/config/pool -id pool_1 extend -disk vdisk_1,vdisk_5
```

```
Storage system address: 10.0.0.2
Storage system port: 443
HTTPS connection

ID = pool_1
Operation completed successfully.
```

View pools

View a list of pools. You can filter on the pool ID.

 **NOTE:** The `show action command` explains how to change the output format.

Format

```
/stor/config/pool {-id <value> | -name <value>}} show
```

Object qualifiers

Qualifier	Description
-id	Type the ID of a pool.
-name	Type the name of a pool.

Example 1 (physical deployments only)

The following command shows details about all pools on a hybrid system:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                                     = pool_13
      Type                                     = Dynamic
      Name                                     = Dynamic
      Description                             =
      Total space                             = 3815810007040 (3.4T)
      Current allocation                       = 940921012224 (876.3G)
      Preallocated                           = 97118896128 (90.4G)
      Remaining space                         = 2777770098688 (2.5T)
      Subscription                           = 1194021912576 (1.0T)
      Flash percent                           = 39%
      Subscription percent                     = 31%
      Alert threshold                         = 70%
      Drives                                  = 6 x 600.0G SAS; 6 x 400.0G SAS
Flash 3
      Number of drives                        = 12
      RAID level                              = 5
      Stripe length                           = 5
      Rebalancing                             = no
      Rebalancing progress                     =
      Health state                            = OK (5)
      Health details                          = "The component is operating
normally. No action is required."
      FAST Cache enabled                      = yes
      Protection size used                    = 21474836480 (20.0G)
      Non-base size used                      = 21474836480 (20.0G)
      Auto-delete state                       = Idle
      Auto-delete paused                      = no
      Auto-delete pool full threshold enabled = yes
      Auto-delete pool full high water mark   = 95%
      Auto-delete pool full low water mark    = 85%
      Auto-delete snapshot space used threshold enabled = no
      Auto-delete snapshot space used high water mark = 25%
      Auto-delete snapshot space used low water mark = 20%
      Data Reduction space saved               = 18790481920 (17.5G)
      Data Reduction Percent                   = 2%
      Data Reduction Ratio                     = 1.0:1
      All flash pool                           = no

2:      ID                                     = pool_14
      Type                                     = Traditional
      Name                                     = Traditional
      Description                             =
      Total space                             = 15750450380800 (14.3T)
      Current allocation                       = 3999290327040 (3.6T)
      Preallocated                           = 2376320188416 (2.1T)
      Remaining space                         = 9374839865344 (8.5T)
      Subscription                           = 14190303510528 (12.9T)
      Flash percent                           = 0%
      Subscription percent                     = 90%
      Alert threshold                         = 70%
      Drives                                  = 5 x 4.0T NL-SAS
      Number of drives                        = 5
      RAID level                              = 5
      Stripe length                           = 5
      Rebalancing                             = no
      Rebalancing progress                     =
      Health state                            = OK (5)
      Health details                          = "The component is operating
normally. No action is required."
      FAST Cache enabled                      = yes
      Protection size used                    = 0
      Non-base size used                      = 0
```

```

Auto-delete state           = Idle
Auto-delete paused          = no
Auto-delete pool full threshold enabled = yes
Auto-delete pool full high water mark   = 95%
Auto-delete pool full low water mark    = 85%
Auto-delete snapshot space used threshold enabled = no
Auto-delete snapshot space used high water mark = 25%
Auto-delete snapshot space used low water mark = 20%
Data Reduction space saved = 0
Data Reduction Percent     = 0%
Data Reduction Ratio       = 1.0:1
All flash pool             = no

```

Example 2 (physical deployments only)

The following example shows all pools for a model that supports dynamic pools.

```
uemcli -d 10.0.0.2 -u Local/joe -p MyPassword456! /stor/config/pool show -detail
```

```

[Response]
Storage system address: 10.0.0.2
Storage system port: 443
HTTPS connection

1:      ID                               = pool_3
      Type                               = Traditional
      Name                               = MyPool
      Description                         = traditional pool
      Total space                         = 14177955479552 (12.8T)
      Current allocation                  = 0
      Preallocated space                  = 38310387712 (35.6G)
      Remaining space                    = 14177955479552 (12.8T)
      Subscription                        = 0
      Subscription percent                = 0%
      Alert threshold                    = 70%
      Drives                             = 9 x 1.6T SAS Flash 3; 5 x
400.0G SAS Flash 2
      Number of drives                   = 14
      RAID level                         = 5
      Stripe length                      = Mixed
      Rebalancing                        = no
      Rebalancing progress               =
      Health state                       = OK (5)
      Health details                     = "The component is operating
normally. No action is required."
      FAST Cache enabled                 = no
      Protection size used               = 0
      Non-base size used                 = 0
      Auto-delete state                  = Idle
      Auto-delete paused                 = no
      Auto-delete pool full threshold enabled = yes
      Auto-delete pool full high water mark   = 95%
      Auto-delete pool full low water mark    = 85%
      Auto-delete snapshot space used threshold enabled = no
      Auto-delete snapshot space used high water mark = 25%
      Auto-delete snapshot space used low water mark = 20%
      Data Reduction space saved          = 0
      Data Reduction percent              = 0%
      Data Reduction ratio                = 1:1
      All flash pool                     = yes

2:      ID                               = pool_4
      Type                               = Dynamic
      Name                               = dynamicPool
      Description                         =
      Total space                         = 1544309178368 (1.4T)
      Current allocation                  = 0
      Preallocated space                  = 38310387712 (35.6G)
      Remaining space                    = 1544309178368 (1.4T)
      Subscription                        = 0
      Subscription percent                = 0%

```

```

Alert threshold                = 70%
Drives                        = 6 x 400.0G SAS Flash 2
Number of drives              = 6
RAID level                    = 5
Stripe length                = 5
Rebalancing                   = no
Rebalancing progress          =
Health state                  = OK (5)
Health details                 = "The component is operating
normally. No action is required."
Protection size used          = 0
Non-base size used            = 0
Auto-delete state             = Idle
Auto-delete paused            = no
Auto-delete pool full threshold enabled = yes
Auto-delete pool full high water mark = 95%
Auto-delete pool full low water mark = 85%
Auto-delete snapshot space used threshold enabled = no
Auto-delete snapshot space used high water mark = 25%
Auto-delete snapshot space used low water mark = 20%
Data Reduction space saved    = 0
Data Reduction percent        = 0%
Data Reduction ratio          = 1:1
All flash pool                = yes

```

Example 3 (virtual deployments only)

The following command shows details for all pools on a virtual system.

```
uemcli -d 10.0.0.3 -u Local/joe -p MyPassword456! /stor/config/pool show -detail
```

```

Storage system address: 10.0.0.3
Storage system port: 443
HTTPS connection

1:      ID                = pool_1
      Name                = Capacity
      Description          =
      Total space          = 4947802324992 (4.5T)
      Current allocation    = 3298534883328 (3T)
      Preallocated space    = 38310387712 (35.6G)
      Remaining space       = 4947802324992 (1.5T)
      Subscription          = 10995116277760 (10T)
      Subscription percent  = 222%
      Alert threshold       = 70%
      Drives                = 1 x 120GB Virtual; 1 x 300GB
Virtual
      Number of drives      = 2
      Health state          = OK (5)
      Health details        = "The component is operating
normally. No action is required."
      Non-base size used    = 1099511625 (1G)
      Auto-delete state     = Running
      Auto-delete paused    = no
      Auto-delete pool full threshold enabled = yes
      Auto-delete pool full high water mark = 95%
      Auto-delete pool full low water mark = 85%
      Auto-delete snapshot space used threshold enabled = yes
      Auto-delete snapshot space used high water mark = 25%
      Auto-delete snapshot space used low water mark = 20%

```

Delete pools

Delete a pool.


Format

```
/stor/config/pool {-id <value> | -name <value>} delete [-async]
```

Object qualifiers

Qualifier	Description
-id	Type the ID of the pool to delete.
-name	Type the name of the pool to delete.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.  NOTE: Simultaneous commands, asynchronous or synchronous, may fail if they conflict in trying to manage the same system elements.

Example

The following deletes pool pool_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool -id pool_1 delete
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Manage pool tiers

Storage tiers allow users to move data between different types of drives in a pool to maximize storage efficiency. Storage tiers are defined by the following characteristics:

- Drive performance.
- Drive capacity.


The following table lists the attributes for storage profiles:

Table 9. Storage tier attributes

Attribute	Description
Name	Storage tier name.
Drives	The list of drive types, and the number of drives of each type in the storage tier.
RAID level (physical deployments only)	RAID level of the storage tier.
Stripe length (physical deployments only)	Comma-separated list of the stripe length of the drives in the storage tier.
Total space	Total capacity in the storage tier.
Current allocation	Currently allocated space.
Remaining space	Remaining space.

View storage tiers

View a list of storage tiers. You can filter on the pool ID.

 **NOTE:** The `show` action command explains how to change the output format.

Format

```
/stor/config/pool/tier {-pool <value> | -poolName <value>} show
```

Object qualifiers

Qualifier	Description
<code>-pool</code>	Type the ID of a pool.
<code>-poolName</code>	Type the name of a pool.

Example 1 (physical deployments only)

The following command shows tier details about the specified pool:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool/tier -pool pool_1 show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:   Name           = Extreme Performance
     Drives        = 6 x 800.0G SAS Flash 2
     Drive type     = SAS Flash
     RAID level     = 5
     Stripe length  = 5
     Total space    = 3111972241408 (2.8T)
     Current allocation = 0
     Remaining space = 3111972241408 (2.8T)
     Spare space configuration = 1

2:   Name           = Performance
     Drives        = 6 x 600.0G SAS
     Drive type     = SAS
     RAID level     = 5
     Stripe length  = 5
     Total space    = 2317671727104 (2.1T)
     Current allocation = 0
     Remaining space = 2317671727104 (2.1T)
     Spare space configuration = 2
```

Example 2 (virtual deployments only)

The following command shows details about the specified pool on a virtual system.

```
uemcli -d 10.0.0.2 -u Local/joe -p MyPassword456! /stor/config/pool/tier -pool pool_1 show -detail
```

```
Storage system address: 10.0.0.2
Storage system port: 443
HTTPS connection

1:   Name           = Extreme Performance
     Drives        =
     Total space    = 0
```

```

Current allocation = 0
Remaining space   = 0

2:  Name           = Performance
    Drives         = 1 x 500GB Virtual
    Total space     = 631242752000 (500.0G)
    Current allocation = 12624855040 (10.0G)
    Remaining space  = 618617896960 (490.0G)

3:  Name           = Capacity
    Drives         =
    Total space     = 0
    Current allocation = 0
    Remaining space  = 0

```

View pool resources

This command displays a list of storage resources allocated in a pool. This can be storage resources provisioned on the specified pool and NAS servers that have file systems allocated in the pool.

The following table lists the attributes for pool resources.

Table 10. Pool resources

Attribute	Description
ID	Storage resource identifier.
Name	Name of the storage resource.
Resource type	Type of the resource. Valid values are: <ul style="list-style-type: none"> LUN File system LUN group VMware NFS VMware VMFS NAS server
Pool	Name of the pool.
Total pool space used	Total space in the pool used by a storage resource. This includes primary data used size, snapshot used size, and metadata size. Space in the pool can be freed if snapshots and thin clones for storage resources are deleted, or have expired.
Total pool space preallocated	Total space reserved from the pool by the storage resource for future needs to make writes more efficient. The pool may be able to reclaim some of this if space is running low. Additional pool space can be freed if snapshots or thin clones are deleted or expire, and also if Data Reduction is applied.
Total pool non-base space used	Total pool space used by snapshots and thin clones.
Health state	Health state of the file system. The health state code appears in parentheses.
Health details	Additional health information. See Appendix A, Reference, for health information details.

Format

```
/stor/config/pool/sr [{-pool <value> | -poolName <value>}] show
```

Object qualifiers

Qualifier	Description
-pool	Type the ID of the pool.
-poolName	Type the name of the pool.

Example

The following command shows details for all storage resources associated with the pool pool_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool/sr -pool pool_1 show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
1:      ID              = res_1
      Name              = File_System_1
      Resource type     = File_System
      Pool              = pool_1
      Total pool space used = 53024473088 (49.3G)
      Total pool preallocated = 15695003648 (14.6G)
      Total pool snapshot space used = 7179124736 (6.6G)
      Total pool non-base space used = 7179124736 (6.6G)
      Health state       = OK (5)
      Health details     = "The component is operating normally. No
action is required."

2:      ID              = sv_1
      Name              = AF_LUN 1
      Resource type     = LUN
      Pool              = pool_1
      Total pool space used = 14448566272 (13.4G)
      Total pool preallocated = 4610351104 (4.2G)
      Total pool snapshot space used = 4593991680 (4.2G)
      Total pool non-base space used = 4593991680 (4.2G)
      Health state       = OK (5)
      Health details     = "The LUN is operating normally. No action is
required."

3:      ID              = res_2
      Name              = File_System_2
      Resource type     = File_System
      Pool              = pool_1
      Total pool space used = 117361025024 (109.3G)
      Total pool preallocated = 3166494720 (2.9G)
      Total pool snapshot space used = 41022308352 (38.2G)
      Total pool non-base space used = 41022308352 (38.2G)
      Health state       = OK (5)
      Health details     = "The component is operating normally. No
action is required."

4:      ID              = sv_2
      Name              = AF_LUN 2
      Resource type     = LUN
      Pool              = pool_1
      Total pool space used = 9500246016 (8.8G)
      Total pool preallocated = 2579349504 (2.4G)
      Total pool snapshot space used = 0
      Total pool non-base space used = 0
      Health state       = OK (5)
      Health details     = "The LUN is operating normally. No action is
required."

5:      ID              = res_3
      Name              = CG1
```

```

Resource type      = LUN group
Pool               = pool_1
Total pool space used      = 892542287872 (831.2G)
Total pool preallocated    = 8863973376 (8.2G)
Total pool snapshot space used = 231799308288 (215.8G)
Total pool non-base space used = 231799308288 (215.8G)
Health state         = OK (5)
Health details       = "The component is operating normally. No
action is required."

```

View storage profiles (physical deployments only)

Storage profiles are preconfigured settings for configuring pools based on the following:

- Types of storage resources that will use the pools.
- Intended usage of the pool.

For example, create a pool for file system storage resources intended for general use. When configuring a pool, specify the ID of the storage profile to apply to the pool.

NOTE: Storage profiles are not restrictive with regard to storage provisioning. For example, you can provision file systems from an FC or iSCSI database pool. However, the characteristics of the storage will be best suited to the indicated storage resource type and use.

Each storage profile is identified by an ID.

The following table lists the attributes for storage profiles.

Table 11. Storage profile attributes

Attribute	Description
ID	ID of the storage profile.
Type	(Available only for systems that support dynamic pools) Type of pool the profile can create. Value is one of the following: <ul style="list-style-type: none"> • Dynamic • Traditional
Description	Brief description of the storage profile.
Drive type	Types of drives for the storage profile.
RAID level	RAID level number for the storage profile. Value is one of the following: <ul style="list-style-type: none"> • 1 - RAID level 1. • 5 - RAID level 5. • 6 - RAID level 6. • 10 - RAID level 1+0.
Maximum capacity	Maximum storage capacity for the storage profile.
Stripe length	Number of drives the data is striped across. NOTE: For best fit profiles, this value is Best fit.
Disk group	List of drive groups recommended for the storage pool configurations of the specified storage profile. This is calculated only when the <code>-configurable</code> option is specified.
Maximum drives to configure	List of the maximum number of drives allowed for the specified storage profile in the recommended drive groups. This is calculated only when the <code>-configurable</code> option is specified.
Maximum capacity to configure	List of the maximum number of free capacity of the drives available to configure for the storage profile in the recommended drive groups. This is calculated only when the <code>-configurable</code> option is specified.

NOTE: The `show action command` explains how to change the output format.

Format

```
/stor/config/profile [-id <value> | -driveType <value> [-raidLevel <value>] | -traditional]
[-configurable] show
```

Object qualifier

Qualifier	Description
-id	Type the ID of a storage profile.
-driveType	Specify the type of drive.
-raidLevel	Specify the RAID type of the profile.
-traditional	(Available only for systems that support dynamic pools) Specify this option to view the profiles that you can use for creating traditional pools. To view the profiles you can use for creating dynamic pools, omit this option.
-configurable	Show only profiles that can be configured, that is, those with non-empty drive group information. If specified, calculates the following drive group information for each profile: <ul style="list-style-type: none">• Disk group• Maximum drives to configure• Maximum capacity to configure If the profile is for a dynamic pool, the calculated information indicates whether the drive group has enough drives for pool creation. The calculation assumes that the pool will be created with the drives in the specified drive group only.

Example 1

The following command shows details for storage profiles that can be used to create dynamic pools:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/profile -configurable show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:  ID              = profile_22
    Type            = Dynamic
    Description     = SAS Flash 2 RAID5 (4+1)
    Drive type      = SAS Flash 2
    RAID level      = 5
    Maximum capacity = 4611148087296 (4.1T)
    Stripe length   = Maximum capacity
    Disk group      =
    Maximum drives to configure =
    Maximum capacity to configure =

2:  ID              = profile_30
    Type            = Dynamic
    Description     = SAS Flash 2 RAID10 (1+1)
    Drive type      = SAS Flash 2
    RAID level      = 10
    Maximum capacity = 9749818597376 (8.8T)
    Stripe length   = 2
    Disk group      =
    Maximum drives to configure =
    Maximum capacity to configure =

3:  ID              = profile_31
    Type            = Dynamic
    Description     = SAS Flash 2 RAID10 (2+2)
    Drive type      = SAS Flash 2
    RAID level      = 10
```

```

Maximum capacity      = 9749818597376 (8.8T)
Stripe length        = 4
Disk group            =
Maximum drives to configure =
Maximum capacity to configure =

```

Example 2

The following command shows details for storage profiles that can be used to create traditional pools in models that support dynamic pools:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/profile -traditional -configurable show
```

```

Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:  ID              = tprofile_22
    Type            = Traditional
    Description      = SAS Flash 3 RAID5 (4+1)
    Drive type       = SAS Flash 3
    RAID level       = 5
    Maximum capacity = 4611148087296 (4.1T)
    Stripe length    = Maximum capacity
    Disk group       = dg_16
    Maximum drives to configure = 5
    Maximum capacity to configure = 1884243623936 (1.7T)

2:  ID              = tprofile_30
    Type            = Traditional
    Description      = SAS Flash 3 RAID10 (1+1)
    Drive type       = SAS Flash 3
    RAID level       = 10
    Maximum capacity = 9749818597376 (8.8T)
    Stripe length    = 2
    Disk group       = dg_13, dg_15
    Maximum drives to configure = 10, 10
    Maximum capacity to configure = 1247522127872 (1.1T), 2954304921600 (2.6T)

3:  ID              = tprofile_31
    Type            = Traditional
    Description      = SAS Flash 3 RAID10 (2+2)
    Drive type       = SAS Flsh 3
    RAID level       = 10
    Maximum capacity = 9749818597376 (8.8T)
    Stripe length    = 4
    Disk group       = dg_13, dg_15
    Maximum drives to configure = 8, 8
    Maximum capacity to configure = 2363443937280 (2.1T), 952103075840 (886.7G)

```

Manage drive groups (physical deployments only)

Drive groups are the groups of drives on the system with similar characteristics, including type, capacity, and spindle speed. When configuring pools, you select the drive group to use and the number of drives from the group to add to the pool.

Each drive group is identified by an ID.

The following table lists the attributes for drive groups.

Table 12. Drive group attributes

Attribute	Description
ID	ID of the drive group.
Drive type	Type of drives in the drive group.

Table 12. Drive group attributes (continued)

Attribute	Description
FAST Cache	Indicates whether the drive group's drives can be added to FAST Cache.
Drive size	Capacity of one drive in the drive group.
Rotational speed	Rotational speed of the drives in the group.
Number of drives	Total number of drives in the drive group.
Unconfigured drives	Total number of drives in the drive group that are not in a pool.
Capacity	Total capacity of all drives in the drive group.
Recommended number of spares	Number of spares recommended for the drive group.
Drives past EOL	Number of drives past EOL (End of Life) in the group.
Drives approaching EOL	Number of drives that will reach EOL in 0-30 days, 0-60 days, 0-90 days and 0-180 days.

View drive groups

View details about drive groups on the system. You can filter on the drive group ID.

 **NOTE:** The `show` action command explains how to change the output format.

Format

```
/stor/config/dg [-id <value>] [-traditional] show
```

Object qualifier

Qualifier	Description
-id	Type the ID of a drive group.
-traditional	(Available only for systems that support dynamic pools) Specify this qualifier to have the system assume that the pools to be created are traditional pools.

Example 1

The following command shows details about all drive groups that can be used to configure dynamic pools:

 **NOTE:** For dynamic pools, the Unconfigured drives value does not include HDD system drives.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/dg show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                      = dg_11
      Drive type                = SAS
      FAST Cache                = no
      Drive size                = 1181778411520 (1.0T)
      Vendor size               = 1.2T
      Rotational speed          = 10000 rpm
      Number of drives          = 44
      Unconfigured drives       = 40
```

```

Capacity = 51998250106880 (47.2T)
Recommended number of spares = 0
Drives past EOL = 0
Drives approaching EOL = 0 (0-30 days), 0 (0-60 days), 0 (0-90 days), 0
(0-180 days)

2: ID = dg_40
Drive type = SAS Flash 4
FAST Cache = no
Drive size = 1881195675648 (1.7T)
Vendor size = 1.9T
Rotational speed = 0 rpm
Number of drives = 18
Unconfigured drives = 18
Capacity = 33861522161664 (30.7T)
Recommended number of spares = 0
Drives past EOL = 0
Drives approaching EOL = 1 (0-30 days), 2 (0-60 days), 2 (0-90 days), 3
(0-180 days)

```

Example 2

The following command shows details about all drive groups that can be used to configure traditional pools in models that support dynamic pools:

NOTE: For dynamic pools, the Unconfigured drives value does not include HDD system drives. For traditional pools, the Unconfigured drives value includes HDD system drives.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/dg -traditional show
```

```

Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1: ID = dg_11
Drive type = SAS
FAST Cache = no
Drive size = 1181778411520 (1.0T)
Vendor size = 1.2T
Rotational speed = 10000 rpm
Number of drives = 44
Unconfigured drives = 44
Capacity = 51998250106880 (47.2T)
Recommended number of spares = 2

2: ID = dg_40
Drive type = SAS Flash 4
FAST Cache = no
Drive size = 1881195675648 (1.7T)
Vendor size = 1.9T
Rotational speed = 0 rpm
Number of drives = 18
Unconfigured drives = 18
Capacity = 33861522161664 (30.7T)
Recommended number of spares = 1

```

View recommended drive group configurations

View the recommended drive groups from which to add drives to a pool based on a specified storage profile or pool type.

NOTE: The `show action command` explains how to change the output format.

Format

```
/stor/config/dg recom {-profile <value>| -pool <value> | -poolName <value>}
```


Action qualifier

Qualifier	Description
-profile	Type the ID of a storage profile. The output will include the list of drive groups recommended for the specified storage profile.
-pool	Type the ID of a pool. The output will include the list of drive groups recommended for the specified pool.
-poolName	Type the name of a pool. The output will include the list of drive groups recommended for the specified pool.

Example 1

The following command shows the recommended drive groups for pool pool_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/dg recom -pool pool_1
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                = DG_1
      Drive type          = SAS
      Drive size          = 536870912000 (500GB)
      Number of drives    = 8
      Allowed numbers of drives = 4,8
      Capacity            = 4398046511104 (4TB)

2:      ID                = DG_2
      Drive type          = SAS
      Drive size          = 268435456000 (250GB)
      Number of drives    = 4
      Allowed numbers of drives = 4
      Capacity            = 1099511627776 (1TB)
```

Manage storage system capacity settings

The following table lists the general storage system capacity attributes:

Table 13. General storage system capacity attributes




Attributes	Description
Free space	Specifies the amount of space that is free (available to be used) in all storage pools on the storage system.
Used space	Specifies the amount of space that is used in all storage pools on the storage system.
Preallocated space	Space reserved across all of the pools on the storage system. This space is reserved for future needs of storage resources, which can make writes more efficient. Each pool may be able to reclaim preallocated space from storage resources if the storage resources are not using the space, and the pool space is running low.
Total space	Specifies the total amount of space, both free and used, in all storage pools on the storage system.
Data Reduction space saved	Specifies the storage size saved on the entire system when using data reduction.  NOTE: Data reduction is available for thin LUNs and thin file systems. The thin file systems must be created on Unity systems running version 4.2.x or later.
Data Reduction percent	Specifies the storage percentage saved on the entire system when using data reduction.

Table 13. General storage system capacity attributes (continued)

Attributes	Description
	 NOTE: Data reduction is available for thin LUNs and thin file systems. The thin file systems must be created on Unity systems running version 4.2.x or later.
Data Reduction ratio	Specifies the ratio between data without data reduction and data after data reduction savings.  NOTE: Data reduction is available for thin LUNs and thin file systems. The thin file systems must be created on Unity systems running version 4.2.x or later.

View system capacity settings

View the current storage system capacity settings.

Format

```
/stor/general/system show
```

Example

The following command displays details about the storage capacity on the system:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/general/system show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      Free space                = 4947802324992 (1.5T)
        Used space                = 4947802324992 (1.5T)
        Total space              = 9895604649984 (3.0T)
        Preallocated space       = 60505210880 (56.3G)
        Data Reduction space saved = 4947802324992 (1.5T)
        Data Reduction percent    = 50%
        Data Reduction ratio      = 1
```

Manage system tier capacity settings

The following table lists the general system tier capacity attributes:

Table 14. General system tier capacity attributes

Attributes	Description
Name	Name of the tier. One of the following: <ul style="list-style-type: none">• Extreme Performance• Performance• Capacity
Free space	Specifies the amount of space that is free (available to be used) in the tier.
Used space	Specifies the amount of space that is used in the tier.
Total space	Specifies the total amount of space, both free and used, in the tier.

View system tier capacity

View the current system tier capacity settings.

Format

```
/stor/general/tier show
```

Example

The following command displays details about the storage tier capacity on the system:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/general/tier show -detail
```

```
1:      Name      = Extreme Performance Tier
      Free space  = 4947802324992 (1.5T)
      Used space  = 4947802324992 (1.5T)
      Total space = 9895604649984 (3.0T)

2:      Name      = Performance Tier
      Free space  = 0
      Used space  = 0
      Total space = 0

3:      Name      = Capacity Tier
      Free space  = 4947802324992 (1.5T)
      Used space  = 4947802324992 (1.5T)
      Total space = 9895604649984 (3.0T)
```

Change disk settings (virtual deployments only)

Change settings of an existing disk.

Format


```
/env/disk -id <value> set [-async] [-name <value>] [-tier <value>]
```

Object qualifier

Qualifier	Description
-id	Disk identifier.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.
-name	Specify the new name for the disk.
-tier	Specify the new tier. Valid values are: <ul style="list-style-type: none">• capacity• performance• extreme

Qualifier	Description
	 NOTE: Disks without a tier cannot be used for pool provisioning.

Example

The following command changes the name of the virtual disk with the ID "vdisk_1".

```
uemcli -d 10.0.0.2 -u Local/joe -p MyPassword456! /env/disk -id vdisk_1 set -name "High-performance storage"
```

```
Storage system address: 10.0.0.2
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Pool use cases

This section describes different CLI use cases for pools.

Create a pool using drives with specific characteristics

This example applies to hybrid Flash arrays, which only support traditional pools.

Retrieve the list of storage profiles

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/profile -configurable show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID              = profile_22
      Description      = SAS RAID5
      Drive type       = SAS
      RAID level       = 5
      Maximum capacity = 4611148087296 (4.1T)
      Stripe length    = Maximum capacity
      Disk group       = dg_16
      Maximum drives to configure = 5
      Maximum capacity to configure = 1884243623936 (1.7T)

2:      ID              = profile_30
      Description      = SAS RAID10 (1+1)
      Drive type       = SAS
      RAID level       = 10
      Maximum capacity = 9749818597376 (8.8T)
      Stripe length    = 2
      Disk group       = dg_13, dg_15
      Maximum drives to configure = 10, 10
      Maximum capacity to configure = 1247522127872 (1.1T), 2954304921600 (2.6T)

3:      ID              = profile_31
      Description      = SAS RAID10 (2+2)
      Drive type       = SAS
      RAID level       = 10
      Maximum capacity = 9749818597376 (8.8T)
      Stripe length    = 4
      Disk group       = dg_13, dg_15
```

```
Maximum drives to configure      = 8, 8
Maximum capacity to configure    = 2363443937280 (2.1T), 952103075840 (886.7G)
```

Configure a new pool

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name MyPool
-description "My custom pool" -storProfile profile_22 -diskGroup dg_16 -drivesNumber 5
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = GP_4
Operation completed successfully.
```

Configure a dynamic pool

You can configure dynamic pools for all-Flash models of Unity running OE version 4.2.x or later. New pools created for these models are dynamic pools by default. Dynamic pools implement advanced RAID technology. In dynamic pools, a RAID group is spread across drive extents in multiple drives. The required spare space is also spread across drive extents in multiple drives. When a drive fails, the extents on the failed drive are rebuilt to spare space extents within the pool.

When you configure dynamic pools, you can select different capacity drives from different drive groups with the same Flash drive types to create a tier. The total drive count of the drive type must be at least the stripe width plus one. For example, the total drive count for a RAID 4 + 1 group must be at least 6.

Step 1: View the list of available drive groups

View the list of available drive groups, as shown in the following example:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/dg show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:  ID                      = dg_2
    Drive type              = SAS Flash 2
    FAST Cache              = yes
    Drive size              = 196971960832 (183.4G)
    Vendor size             = 200.0G
    Rotational speed        = 0 rpm
    Number of drives        = 12
    Unconfigured drives     = 12
    Capacity                = 2363663529984 (2.1T)
    Recommended number of spares = 0

2:  ID                      = dg_3
    Drive type              = SAS Flash 2
    FAST Cache              = yes
    Drive size              = 393846128640 (366.7G)
    Vendor size             = 400.0G
    Rotational speed        = 0 rpm
    Number of drives        = 12
    Unconfigured drives     = 6
    Capacity                = 4726153543680 (4.2T)
    Recommended number of spares = 0
```

Step 2: View the list of storage profiles

View the list of storage profiles, as shown in the following example:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/profile show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                      = profile_1
      Type                     = Dynamic
      Description               = SAS Flash 2 RAID5 (4+1)
      Drive type                = SAS Flash 2
      RAID level                = 5
      Maximum capacity          = 97373737844736 (88.5T)
      Stripe length             = 5
      Disk group                =
      Maximum drives to configure =
      Maximum capacity to configure =
```

Step 3: Configure the dynamic pool

Configure the dynamic pool with the specified drive groups and profiles. Optionally set **-type** to **dynamic** and make sure that the **-drivesNumber** value for each drive type is not less than the drive group's RAID group width plus one:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name mypool
-diskGroup dg_2,dg_3 -drivesNumber 4,2 -storProfile profile_1 -type dynamic
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = pool_13
Operation completed successfully.
```

Configure a traditional pool for an all-Flash model

You can configure traditional pools for all-Flash models of Unity running OE version 4.2.x and later by explicitly setting the **Type** attribute to **traditional**. If you do not set **Type** to **traditional** when you create a pool in the Unisphere CLI, a dynamic pool is created.

Step 1: View the list of storage profiles

View the list of storage profiles, as shown in the following example:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/profile -traditional
-configurable show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                      = tprofile_2
      Type                     = Traditional
      Description               = SAS Flash 2 RAID5 (8+1)
      Drive type                = SAS Flash 2
      RAID level                = 5
      Maximum capacity          = 95010661072896 (86.4T)
      Stripe length             = 9
      Disk group                = dg_34, dg_26
      Maximum drives to configure = 9, 9
      Maximum capacity to configure = 60189403250688 (54.7T), 2232208064512 (2.0T)

2:      ID                      = tprofile_4
      Type                     = Traditional
      Description               = SAS Flash 2 RAID5
      Drive type                = SAS Flash 2
      RAID level                = 5
```

```
Maximum capacity          = 95010661072896 (86.4T)
Stripe length            = Maximum capacity
Disk group                = dg_34, dg_26
Maximum drives to configure = 9, 10
Maximum capacity to configure = 60189403250688 (54.7T), 2691354329088 (2.4T)
```

Step 2: Configure the traditional pool

Configure a traditional pool with the specified profile. Make sure you set `-type` to **traditional** and that the `-drivesNumber` is a multiple of the RAID group width.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool create -name test
-diskGroup dg_34 -drivesNumber 9 -storProfile tprofile_2 -type traditional
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = pool_6
Operation completed successfully.
```

Add drives to an existing pool

Retrieve the list of existing pools

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /store/config/pool show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID          = SPL_1
      Name          = Performance
      Description    =
      Free space     = 408944640 (390G)
      Capacity       = 1099511627776 (1T)
      Drives         = 6 x 250GB SAS
      Number of drives = 6
      Unused drives  = 1
      RAID level     = 5
      System pool    = yes

2:      ID          = SPL_2
      Name          = Capacity
      Description    =
      Free space     = 1319413953331 (1.2T)
      Capacity       = 13194139533312 (12T)
      Drives         = 8 x 2GB NL-SAS
      Number of drives = 8
      Unused drives  = 0
      RAID level     = 6
      System pool    = yes

3:      ID          = SPL_3
      Name          = Extreme Performance
      Description    =
      Free space     = 209715200 (200M)
      Capacity       = 322122547200 (300G)
      Drive type     = EFD
      Number of drives = 4
      Unused drives  = 0
      RAID level     = 5
      System pool    = yes
```

Retrieve the list of recommended disk groups for the selected pool

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/dg recom -pool SPL_3
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                      = DG_4
        Drive type              = EFD
        Drive size              = 107374182400 (100G)
        Number of drives        = 4
        Allowed numbers of drives = 4
        Capacity                = 419430400 (400G)
```

Extend the existing pool

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool -id SPL_3 extend -
diskGroup DG_4 -drivesNumber 4
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

ID = SPL_3
Operation completed successfully.
```

The show action command

The `show` action command displays a list of objects that exist on the system and the attributes of those objects. You can specify an object qualifier to view the attributes for a single object. The `show` action command provides qualifiers for changing the display of the output, including the format and the attributes to include. The available output formats are name-value pair (NVP), table, and comma-separated values (CSV).

Format

```
uemcli [<switches>] <object> [<object qualifier>] show [{-detail | -brief | -filter
<value>}] [-output {nvp | table [-wrap] | csv}]
```

Action qualifier

Qualifier	Description
<code>-output -o</code>	Specify the output format. Value is one of the following: <ul style="list-style-type: none"><code>nvp</code> — The name-value pair (NVP) format displays output as <code>name=value</code>. Name-value pair format provides an example of the NVP format.<code>table</code> — The table format displays output as a table, with column headers and rows. By default, values that are too long to fit in a column are cut off. Add <code>-wrap</code> after the table qualifier, separated by a space, so that the values wrap. Table format provides an example of the table format.<code>csv</code> — The comma-separated values (CSV) format is similar to the table format, but the names and values are separated by commas. Comma-separated values format provides an example of the CSV format.
<code>-detail</code>	Display all attributes.
<code>-brief</code>	Display only the basic attributes (default).
<code>-filter</code>	Comma-separated list of attributes which are included into the command output.

Name-value pair format

```
1:      ID          = la0_SPA
      SP           = SPA
      Ports        = eth0_SPA,eth1_SPA
      Health state  = OK (5)

2:      ID          = la0_SPB
      SP           = SPB
      Ports        = eth0_SPB,eth1_SPB
      Health state  = OK (5)
```

Table format

ID	SP	Ports	Health state
la0_SPA	SPA	eth0_SPA,eth1_SPA	OK (5)
la0_SPB	SPB	eth0_SPB,eth1_SPB	OK (5)

Comma-separated values format

```
ID,SP,Ports,Health state
la0_SPA,SPA,"eth0_SPA,eth1_SPA",OK (5)
la0_SPB,SPB,"eth0_SPB,eth1_SPB",OK (5)
```

Example

The following command modifies the set of attributes in the show action output. For example, if you add `-filter "ID,ID,ID,ID"` to the command, in the output you will see four lines with the "ID" attribute for each listed instance:

```
1:  ID = la_0
    ID = la_0
    ID = la_0
    ID = la_0
```

```
uemcli /net/nas/server show -filter "ID, SP, Health state, ID, Name"
```

Filter format

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection
```

```
1:      ID          = nas_1
      SP           = SPA
      Health state  = OK (5)
      ID           = nas_1
      Name          = Mynas1

2:      ID          = nas_2
      SP           = SPA
      Health state  = OK (5)
      ID           = nas_2
      Name          = Mynas2
```

Managing the FAST Cache using the CLI (supported physical deployments only)

Topics:

- [Manage FAST Cache](#)
- [Manage FAST Cache storage objects \(physical deployments only\)](#)

Manage FAST Cache

FAST Cache is a storage efficiency technology that uses disks to expand the cache capability of the storage system to provide improved performance.

The following table lists the attributes for FAST Cache:

Table 15. FAST Cache attributes

Attribute	Description
Capacity	Capacity of the FAST Cache.
Drives	The list of disk types, and the number of disks of each type in the FAST Cache.
Number of drives	Total number of disks in the FAST Cache.
RAID level	RAID level applied to the FAST Cache disks. This value is always RAID 1 .
Health state	Health state of the FAST Cache. The health state code appears in parentheses.
Health details	Additional health information. See Appendix A, Reference, for health information details.


Create FAST Cache

Configure FAST Cache. The storage system generates an error if FAST Cache is already configured.

Format

```
/stor/config/fastcache create [-async] -diskGroup <value> -drivesNumber <value> [-enableOnExistingPools {yes | no}]
```

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.
-diskGroup	Specify the drive group to include in the FAST Cache.  NOTE: Only SAS Flash 2 drives can be used in the FAST Cache.

Qualifier	Description
-drivesNumber	Specify the number of drives to include in the FAST Cache.
-enableOnExistingPools	Specify whether FAST Cache is enabled on all existing pools. Valid values are: <ul style="list-style-type: none"> • yes • no

Example

The following command configures FAST Cache with six drives from drive group dg_2, and enables FAST Cache on existing pools:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastcache create -diskGroup dg_2 -drivesNumber 6 -enableOnExistingPools yes
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

View FAST Cache settings

View the FAST Cache parameters.

Format

```
/stor/config/fastcache show
```

Example

The following command displays the FAST Cache parameters for a medium endurance Flash drive:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastcache show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      Total space           = 536870912000 (500G)
        Drives                = 6 x 200GB SAS Flash 2
        Number of drives      = 6
        RAID level            = 1
        Health state          = OK (5)
        Health details         = "The component is operating normally.  No action is
required."
```

Extend FAST Cache

Extend the FAST Cache by adding more drives.

Format

```
/stor/config/fastcache extend [-async] -diskGroup <value> -drivesNumber <value>
```

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.
-diskGroup	Specify the comma-separated list of SAS Flash drives to add to the FAST Cache. Any added drives must have the same drive type and drive size as the existing drives.
-drivesNumber	Specify the number of drives for each corresponding drive group to be added to the FAST Cache.

Example

The following command adds six drives from drive group "dg_2" to FAST cache.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastcache extend -diskGroup dg_2 -drivesNumber 6
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Shrink FAST Cache

Shrink the FAST Cache by removing storage objects.

Format

```
/stor/config/fastcache shrink [-async] -so <value>
```

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.
-so	Specify the comma-separated list of storage objects to remove from the FAST Cache. Run the <code>/stor/config/fastcache/so show</code> command to obtain a list of all storage objects currently in the FAST Cache.

Example

The following command removes Raid Group RG_1 from the FAST Cache.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastcache shrink -so rg_1
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Delete FAST Cache

Delete the FAST Cache configuration. The storage system generates an error if FAST Cache is not configured on the system.

Format

```
/stor/config/fastcache delete [-async]
```

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.

Example

The following command deletes the FAST Cache configuration:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastcache delete
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Manage FAST Cache storage objects (physical deployments only)

FAST Cache storage objects include the RAID groups and drives that are in the FAST Cache.

Table 16. FAST Cache storage object attributes

Attribute	Description
ID	Identifier of the storage object.
Type	Type of storage object.
RAID level	RAID level applied to the storage object.
Drive type	Type of drive.
Number of drives	Number of drives in the storage object.
Drives	Comma-separated list of the drive IDs for each storage object.
Total space	Total space used by the storage object.
Device state	The status of the FAST Cache device. Values are: <ul style="list-style-type: none">• OK - This cache device is operating normally.• Degraded - One drive of this cache device is faulted.• Faulted - This cache device cannot operate normally.• Expanding - This cache device is expanding.• Expansion Ready - This cache device finished expanding.• Expansion Failure - This cache device failed to expand.• Shrinking - This cache device is shrinking.• Shrink Done - This cache device has flushed pages and is removed from FAST Cache.

View FAST Cache storage objects

View a list of all storage objects, including RAID groups and drives, that are in the FAST Cache.

Format

```
/stor/config/fastcache/so [-id <value> ] show
```

Object qualifier

Qualifier	Description
-id	Type the ID of the storage object in the FAST Cache.

Example 1

The following example shows FAST Cache storage objects on the system.

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastcache/so show
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1:      ID                = rg_6
        Type              = RAID group
        Stripe length     = 2
        RAID level        = 1
        Number of drives  = 2
        Drive type        = SAS Flash 2
        Drives            = dae_0_1_disk_1, dae_0_1_disk_2
        Total space       = 195400433664 (181.9G)
        Device state      = OK
```

Managing FAST VP using the CLI (supported physical deployments only)

Topics:

- [Manage FAST VP pool settings](#)
- [Manage FAST VP general settings](#)

Manage FAST VP pool settings


Fully Automated Storage Tiering for Virtual Pools (FAST VP) is a storage efficiency technology that automatically moves data between storage tiers within a pool based on data access patterns.

The following table lists the attributes for FAST VP pool settings.

Table 17. FAST VP pool attributes

Attribute	Description
Pool	Identifies the pool.
Status	Identifies the status of data relocation on the pool. Value is one of the following: <ul style="list-style-type: none"> • Not started - Data relocation has not started. • Paused - Data relocation is paused. • Completed - Data relocation is complete. • Stopped by user - Data relocation was stopped by the user. • Active - Data relocation is in progress. • Failed - Data relocation failed.
Relocation type	Type of data relocation. Value is one of the following: <ul style="list-style-type: none"> • Manual - Data relocation was initiated by the user. • Scheduled or rebalancing - Data relocation was initiated by the system because it was scheduled, or because the system rebalanced the data.
Schedule enabled	Identifies whether the pool is rebalanced according to the system FAST VP schedule. Value is one of the following: <ul style="list-style-type: none"> • yes • no
Start time	Indicates the time the current data relocation started.
End time	Indicates the time the current data relocation is scheduled to end.
Data relocated	The amount of data relocated during an ongoing relocation, or the previous relocation if a data relocation is not occurring. The format is: <value> [suffix] where: <ul style="list-style-type: none"> • value - Identifies the size of the data relocated. • suffix - Identifies that the value relates to the previous relocation session.
Rate	Identifies the transfer rate for the data relocation. Value is one of the following: <ul style="list-style-type: none"> • Low - Least impact on system performance. • Medium - Moderate impact on system performance. • High - Most impact on system performance.

Table 17. FAST VP pool attributes (continued)

Attribute	Description
	Default value is medium.  NOTE: This field is blank if data relocation is not in progress.
Data to move up	The amount of data in the pool scheduled to be moved to a higher storage tier.
Data to move down	The amount of data in the pool scheduled to be moved to a lower storage tier.
Data to move within	The amount of data in the pool scheduled to be moved within the same storage tiers for rebalancing.
Data to move up per tier	The amount of data per tier that is scheduled to be moved to a higher tier. The format is: <tier_name>:[value] where: <ul style="list-style-type: none"> tier_name - Identifies the storage tier. value - Identifies the amount of data in that tier to be move up.
Data to move down per tier	The amount of data per tier that is scheduled to be moved to a lower tier. The format is: <tier_name>:[value] where: <ul style="list-style-type: none"> tier_name - Identifies the storage tier. value - Identifies the amount of data in that tier to be moved down.
Data to move within per tier	The amount of data per tier that is scheduled to be moved to within the same tier for rebalancing. The format is: <tier_name>:[value] where: <ul style="list-style-type: none"> tier_name - Identifies the storage tier. value - Identifies the amount of data in that tier to be rebalanced.
Estimated relocation time	Identifies the estimated time required to perform the next data relocation.

View FAST VP pool settings

View FAST VP settings on a pool.

Format

```
/stor/config/pool/fastvp [{-pool <value> | -poolName <value>}] show
```

Object qualifiers

Qualifier	Description
-pool	Type the ID of the pool.
-poolName	Type the name of the pool.

Example

The following command lists the FAST VP settings on the storage system:


```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool/fastvp -show -detail
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

1: Pool                                = pool_1
   Relocation type                     = manual
   Status                             = Active
   Schedule enabled                    = no
   Start time                         = 2013-09-20 12:55:32
   End time                           = 2013-09-20 21:10:17
   Data relocated                     = 100111454324 (100G)
   Rate                               = high
   Data to move up                     = 4947802324992 (4.9T)
   Data to move down                  = 4947802324992 (4.9T)
   Data to move within                 = 4947802324992 (4.9T)
   Data to move up per tier            = Performance: 500182324992 (500G), Capacity:
1000114543245 (1.0T)
   Data to move down per tier          = Extreme Performance: 1000114543245 (1.0T),
Performance: 500182324992 (500G)
   Data to move within per tier        = Extreme Performance: 500182324992 (500G),
Performance: 500182324992 (500G), Capacity: 500182324992 (500G)
   Estimated relocation time          = 7h 30m
```

Change FAST VP pool settings

Modify FAST VP settings on an existing pool.


Format

```
/stor/config/pool/fastvp {-pool <value> | -poolName <value>} set [-async] -schedEnabled
{yes | no}
```

Object qualifiers

Qualifier	Description
-pool	Type the ID of the pool.
-poolName	Type the name of the pool.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.  NOTE: Simultaneous commands, asynchronous or synchronous, may fail if they conflict in trying to manage the same system elements.
-schedEnabled	Specify whether the pool is rebalanced according to the system FAST VP schedule. Value is one of the following: <ul style="list-style-type: none">• yes• no

Example

The following example enables the rebalancing schedule on pool pool_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool/fastvp -pool pool_1 set -schedEnabled yes
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Pool ID = pool_1
Operation completed successfully.
```

Start data relocation

Start data relocation on a pool.


Format

```
/stor/config/pool/fastvp {-pool <value> | -poolName <value>} start [-async] [-rate {low | medium | high}] [-endTime <value>]
```

Object qualifiers

Qualifier	Description
-pool	Type the ID of the pool to resume data relocation.
-poolName	Type the name of the pool to resume data relocation.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.  NOTE: Simultaneous commands, asynchronous or synchronous, may fail if they conflict in trying to manage the same system elements.
-pool	Type the ID of the pool.
-endTime	Specify the time to stop the data relocation. The format is: [HH:MM] where: <ul style="list-style-type: none">• HH — Hour.• MM — Minute. Default value is eight hours from the current time.
-rate	Specify the transfer rate for the data relocation. Value is one of the following: <ul style="list-style-type: none">• Low — Least impact on system performance.• Medium — Moderate impact on system performance.• High — Most impact on system performance. Default value is the value set at the system level.

Example

The following command starts data relocation on pool pool_1, and directs it to end at 04:00:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool/fastvp -pool pool_1
start -endTime 04:00
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Stop data relocation

Stop data relocation on a pool.

Format

```
/stor/config/pool/fastvp {-pool <value> | -poolName <value>} stop [-async]
```

Object qualifiers

Qualifier	Description
-pool	Type the ID of the pool.
-poolName	Type the name of the pool.

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.

Example

The following command stops data relocation on pool pool_1:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/pool/fastvp -pool pool_1
stop
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

Manage FAST VP general settings


Fully Automated Storage Tiering for Virtual Pools (FAST VP) is a storage efficiency technology that automatically moves data between storage tiers within a pool based on data access patterns.

The following table lists the attributes for FAST VP general settings.

Table 18. FAST VP general attributes

Attribute	Description
Paused	Identifies whether the data relocation is paused. Value is one of the following: <ul style="list-style-type: none">yes

Table 18. FAST VP general attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> no
Schedule-enabled	<p>Identifies whether the pool is rebalanced according to the system FAST VP schedule. Value is one of the following:</p> <ul style="list-style-type: none"> yes no
Frequency	<p>Data relocation schedule. The format is:</p> <p>Every <i><days_of_the_week></i> at <i><start_time></i> until <i><end_time></i></p> <p>where:</p> <ul style="list-style-type: none"> <i><days_of_the_week></i> - List of the days of the week that data relocation will run. <i><start_time></i> - Time the data relocation starts. <i><end_time></i> - Time the data relocation finishes.
Rate	<p>Identifies the transfer rate for the data relocation. Value is one of the following:</p> <ul style="list-style-type: none"> Low - Least impact on system performance. Medium - Moderate impact on system performance. High - Most impact on system performance. <p>Default value is medium.</p> <p> NOTE: This field is blank if data relocation is not in progress.</p>
Data to move up	The amount of data in the pool scheduled to be moved to a higher storage tier.
Data to move down	The amount of data in the pool scheduled to be moved to a lower storage tier.
Data to move within	The amount of data in the pool scheduled to be moved within the same storage tiers for rebalancing.
Estimated scheduled relocation time	Identifies the estimated time required to perform the next data relocation.

Change FAST VP general settings

Change FAST VP general settings.

Format

```
/stor/config/fastvp set [-async] [-schedEnabled {yes | no}] [-days <value>] [-at <value>]
[-until <value>] [-rate {low | medium | high}] [-paused {yes | no}]
```

Action qualifier

Qualifier	Description
-async	Run the operation in asynchronous mode.
-paused	<p>Specify whether to pause data relocation on the storage system. Valid values are:</p> <ul style="list-style-type: none"> yes no
-schedEnabled	<p>Specify whether the pool is rebalanced according to the system FAST VP schedule. Valid values are:</p> <ul style="list-style-type: none"> yes no
-days	<p>Specify a comma-separated list of the days of the week to schedule data relocation. Valid values are:</p> <ul style="list-style-type: none"> mon – Monday

Qualifier	Description
	<ul style="list-style-type: none"> • <code>tue</code> – Tuesday • <code>wed</code> – Wednesday • <code>thu</code> – Thursday • <code>fri</code> – Friday • <code>sat</code> – Saturday • <code>sun</code> – Sunday
<code>-at</code>	<p>Specify the time to start the data relocation. The format is:</p> <p>[HH:MM]</p> <p>where:</p> <ul style="list-style-type: none"> • HH – Hour • MM – Minute <p>Valid values are between 00:00 and 23:59. Default value is 00:00.</p>
<code>-until</code>	<p>Specify the time to stop the data relocation. The format is:</p> <p>[HH:MM]</p> <p>where:</p> <ul style="list-style-type: none"> • HH – Hour • MM – Minute <p>Valid values are between 00:00 and 23:59. Default value is eight hours after the time specified with the <code>-at</code> parameter.</p>
<code>-rate</code>	<p>Specify the transfer rate for the data relocation. Value is one of the following:</p> <ul style="list-style-type: none"> • <code>low</code> – Least impact on system performance. • <code>medium</code> – Moderate impact on system performance. • <code>high</code> – Most impact on system performance. <p>Default value is <code>medium</code>.</p>

Example

The following command changes the data relocation schedule to run on Mondays and Fridays from 23:00 to 07:00:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastvp set -schedEnabled yes -days "Mon,Fri" -at 23:00 -until 07:00
```

```
Storage system address: 10.0.0.1
Storage system port: 443
HTTPS connection

Operation completed successfully.
```

View FAST VP general settings

View the FAST VP general settings.

Format

```
/stor/config/fastvp show -detail
```

Example

The following command displays the FAST VP general settings:

```
uemcli -d 10.0.0.1 -u Local/joe -p MyPassword456! /stor/config/fastvp show -detail
```

```
Storage system address: 10.0.0.1  
Storage system port: 443  
HTTPS connection
```

```
1: Paused = no  
Schedule enabled = yes  
Frequency = Every Mon, Fri at 22:30 until 8:00  
Rate = high  
Data to move up = 4947802324992 (1.5T)  
Data to move down = 4947802324992 (1.5T)  
Data to move within = 4947802324992 (1.5T)  
Estimated scheduled relocation time = 7h 30m
```