

# Dell EMC Unisphere for PowerMax Product Guide

Version 9.1.0

## Unisphere for PowerMax Product Guide

September 2019

• Overview of Unisphere.....	3
• Capacity information.....	4
• Login authentication.....	5
• Unisphere dashboards overview.....	5
• Understanding the system health score.....	8
• Managing jobs.....	10
• Unisphere for PowerMax alert monitoring recommendations.....	10
• Server alerts.....	11
• Understanding settings.....	11
• Understanding licenses.....	12
• Understanding user authorization.....	12
• Individual and group roles.....	13
• Roles.....	13
• User IDs.....	14
• Roles and associated permissions.....	15
• RBAC roles for TimeFinder SnapVX local and remote replication actions.....	18
• RBAC roles for SRDF local and remote replication actions.....	18
• Understanding access controls for volumes.....	19
• Understanding Storage Management.....	19
• Understanding storage provisioning.....	21
• Understanding storage groups.....	23
• Understanding compression.....	24
• Understanding service levels.....	24
• Suitability Check restrictions.....	24
• Understanding storage templates.....	25
• Understanding Storage Resource Pools.....	25
• Understanding volumes.....	25
• Understanding Federated Tiered Storage.....	25
• Understanding FAST.....	26
• Understanding Workload Planner.....	26
• Understanding time windows.....	26
• Understanding FAST.X.....	27
• Virtualize external LUNs.....	27
• Understanding tiers.....	28
• Understanding thin pools.....	28

- Understanding disk groups..... 29
- Understanding Virtual LUN Migration..... 29
- Understanding vVols..... 30
- Understanding Host Management..... 30
- Understanding hosts..... 30
- Understanding masking views..... 30
- Understanding port groups..... 31
- Understanding initiators..... 31
- Understanding PowerPath hosts..... 31
- Understanding mainframe management..... 31
- Understanding Data Protection Management..... 32
- Managing remote replication sessions..... 33
- Understanding non-disruptive migration..... 33
- Understanding Virtual Witness..... 35
- Understanding SRDF Delta Set Extension (DSE) pools..... 35
- Understanding TimeFinder/Snap operations..... 35
- Understanding Open Replicator..... 36
- Open Replicator session options..... 36
- Understanding device groups..... 38
- Understanding SRDF groups..... 39
- SRDF session modes..... 39
- SRDF session options..... 40
- SRDF/A control actions..... 43
- SRDF group modes..... 44
- RDFA flags..... 45
- Understanding TimeFinder/Clone operations..... 45
- Understanding TimeFinder/Mirror sessions..... 46
- Understanding TimeFinder SnapVX..... 46
- Understanding RecoverPoint..... 46
- Understanding Performance Management..... 46
- Understanding Database Storage Analyzer (DSA)..... 47
- Understanding Unisphere support for VMware..... 48
- Understanding eNAS..... 48
- Understanding iSCSI..... 48
- Understanding dynamic cache partitioning..... 49

# Overview of Unisphere

Unisphere enables you to configure and manage PowerMax, VMAX All Flash, and VMAX storage systems.

Unisphere is a HTML5 web-based application that enables you to configure and manage PowerMax, VMAX All Flash, and VMAX storage systems. The term Unisphere incorporates "Unisphere for PowerMax" for the management of PowerMax and All Flash storage systems running PowerMaxOS 5978, and "Unisphere for VMAX" for the management of VMAX All Flash and VMAX storage systems running HYPERMAX OS 5977 and Enginuity OS 5876.

Blog posts and videos that provide an overview of Unisphere functionality can be accessed by clicking [here](#).



A HTML5 based Unisphere provides several advantages:

- Improved security
- Modern user interface
- Reduced application response times

The side panel has the following items when the **All Systems** view is selected:

- **HOME**—View system view dashboard of all storage systems being managed
- **PERFORMANCE**—Monitors and manages storage system performance data (Dashboards, Charts, Analyze, Heatmap, Reports, Plan, Real-Time traces, and Performance Database management). Refer to [Understanding Performance Management](#) on page 46 for more information.
- **VMWARE**—Views all the relevant storage-related objects at an ESXi server and helps troubleshooting storage performance-related issues at the ESXi server. Refer to [Understanding Unisphere support for VMware](#) on page 48 for more information.
- **DATABASES**—Monitors and troubleshoots database performance issues. Refer to [Understanding Database Storage Analyzer \(DSA\)](#) on page 47 for more information.
- **EVENTS**—Includes Alerts and Job List.
- **SUPPORT**—Displays support information.

You can set preferences by clicking .

You can hide the side panel by clicking  and you can display the display panel by clicking  again.

You can return to the **All Systems** view by clicking **HOME**.

The side panel has the following items when the storage system-specific view is selected:

- **HOME**—View system view dashboard of all storage systems being managed
- **DASHBOARD**—View the following dashboards for a selected storage system: Capacity and Performance, System Health, Storage Group compliance, Capacity, and Replication
- **STORAGE**—Manage storage (storage groups, service levels, templates, Storage resource Pools, volumes, external storage, vVols, FAST policies, tiers, thin pools, disk groups, and VLUN migration). Refer to [Understanding Storage Management](#) on page 19 for more information.
- **HOSTS**—Manage hosts (hosts, masking views, port groups, initiators, XtremSW Cache Adapters, PowerPath Hosts, mainframe, and CU images). Refer to [Understanding Host Management](#) on page 30 for more information.
- **DATA PROTECTION**—Manage data protection (storage groups, device groups, SRDF groups, migrations, virtual witness, Open Replicator, SRDF/A DSE pools, TimeFinder SnapVX pools,

and RecoverPoint systems). Refer to [Understanding Data Protection Management](#) on page 32 for more information.

- **PERFORMANCE**—Monitors and manages storage system performance data (Dashboards, Charts, Analyze, Heatmap, Reports, Plan, Real-Time traces, and Performance Database management). Refer to [Understanding Performance Management](#) on page 46 for more information.
- **SYSTEM**—Includes Hardware, Properties, File (eNAS), and iSCSI.
- **EVENTS**—Includes Alerts, Job List, and Audit log.
- **SUPPORT**—Displays support information.

The following options are available from the title bar:

- Discover systems
- Refresh system information.
- Search for objects
- View newly added features.
- View and manage alerts.
- View and manage jobs.
- View online help
- Exit the console.

A Unisphere Representational State Transfer (REST) API is also available. The API enables you to access diagnostic, performance and configuration data, and also enables you to perform provisioning operations on the storage system.

#### *Supporting documentation*

Information on the installation of Unisphere for PowerMax can be found in the *Unisphere for PowerMax Installation Guide* at the [Dell EMC support website](#) or the [technical documentation page](#).

For information specific to this Unisphere product release, see the *Unisphere for PowerMax Release Notes* at the [Dell EMC support website](#) or the [technical documentation page](#).

Your comments—Your suggestions help continue to improve the accuracy, organization, and overall quality of the user publications. Send your feedback to: [content feedback](#).

## Capacity information

Unisphere supports measurement of capacity using both the base 2 (binary) and base 10 (decimal) systems.

Storage capacity can be measured using two different systems – base 2 (binary) and base 10 (decimal). Organizations such as the International System of Units (SI) recommend using the base 10 measurement to describe [storage capacity](#). In base 10 notation, one MB is equal to 1 million bytes, and one GB is equal to 1 billion bytes.

Operating systems generally measure storage capacity using the base 2 measurement system. Unisphere and Solutions Enabler use the base 2 measurement system to display storage capacity with the TB notation as it is more universally understood. In base 2 notation, one MB is equal to 1,048,576 bytes and one GB is equal to 1,073,741,824 bytes.

Name	Abbreviation	Binary Power	Binary Value (in Decimal)	Decimal Power	Decimal (Equivalent)
kilobyte	KB	2 <sup>10</sup>	1,024	10 <sup>3</sup>	1,000

Name	Abbreviation	Binary Power	Binary Value (in Decimal)	Decimal Power	Decimal (Equivalent)
megabyte	MB	2 <sup>20</sup>	1,048,576	10 <sup>6</sup>	1,000,000
GB	GB	2 <sup>30</sup>	1,073,741,824	10 <sup>9</sup>	1,000,000,000
terabyte	TB	2 <sup>40</sup>	1,099,511,627,776	10 <sup>12</sup>	1,000,000,000,000

## Login authentication

When you log in, Unisphere checks the following locations for validation:

- **Windows** — The user has a Windows account on the server. (Log in to Unisphere with your Windows Domain\Username and Password.)
- **LDAP-SSL** — The user account is stored on an LDAP-SSL server. (Log in to with your Unisphere LDAP-SSL Username and Password.)  
The Unisphere Administrator or SecurityAdmin must set the LDAP-SSL server location in the LDAP-SSL Configuration dialog box.
- **Local** — The user has a local Unisphere account. Local user accounts are stored locally on the Unisphere server host. (Log in to Unisphere with your Username and Password.)

User names are case-sensitive and allow alphanumeric characters of either case, an underscore, a dash, or a period:

- a-z
- A-Z
- 0-9
- \_
- .
- -

Passwords cannot exceed 16 characters. There are no restrictions on special characters when using passwords.

The Initial Setup User, Administrator, or SecurityAdmin must create a local Unisphere user account for each user.

## Unisphere dashboards overview

Unisphere dashboards display overall status information.

*Home dashboard view for all storage systems*

The home dashboard view for all storage systems (the default view after logging in) provides an overall view of the status of the storage systems that Unisphere manages in terms of the following:

- **Compliance**—Service level compliance data in the form of storage group counts for each compliance state (Critical, Marginal, Stable), total storage group count, and number of storage groups with no service level assigned.
- **Capacity**—Percentage of allocated capacity for the storage system
- **Health score**—an overall health score based on the lowest health score out of the five metrics (see [Understanding the system health score](#) on page 8 for more information).
- **Throughput**—Current throughput for the system, in MB/second

- **IOPS**—Current IOPS for the system
- **Efficiency**—The overall efficiency ratio for the array. It represents the ratio of the sum of all TDEVs plus snapshot sizes (calculated based on the 128K track size) and the physical used storage (calculated based on the compressed pool track size). The ratio is displayed using the k notation to represent 1000, for example, 361900:1 is displayed as 361.9k:1.
- **CloudIQ**—If CLOUDIQ has been selected and the Secure Remote Services (SRS) gateway has not been registered, an option to register is displayed on each storage system card. Clicking **REGISTER** enables you to register the SRS gateway . If the SRS gateway has already been registered, the enabled or disabled data collection status is displayed within each storage system card. Each card also displays the last time data was sent to CloudIQ for that storage system. Clicking the Enabled/Disabled link opens the Settings dialog on the CloudIQ tab to enable the user to enable or disable data collection on the storage system .

Unisphere release 9.1 enables you to tag arrays and storage groups for PowerMaxOS 5978 storage systems running release 9.1. You can add a tag to an object for identification or to give other information. An association between objects can be inferred by adding the same tag to each object. You can search by tag to retrieve related objects. Tags can be added to and removed from objects. When a tag is no longer associated with an object, it no longer exists.

#### *Home dashboard view for a specific storage system*

The home dashboard view for a specific storage system provides a view of the status of a specific storage system. The default view is the view of Performance and Capacity information. The following panels are displayed:

- **SYSTEM HEALTH**
- **SG COMPLIANCE**
- **CAPACITY**
- **REPLICATION**

#### *System Health dashboard view for a specific storage system*

The **System Health** dashboard provides a single place from which you can quickly determine the health of the system. You can also access hardware information.

The System Health section displays values for the following five high-level health or performance metrics: System Utilization, Configuration, Capacity, SG Response Time, and Service Level Compliance. It also displays an overall health score based on these five categories. The overall system health score is based on the lowest health score out of the categories System Utilization, Configuration, Capacity, SG Response Time, and service level compliance. See [Understanding the system health score](#) on page 8 for details on how these scores are calculated. These five categories are for systems running HYPERMAX OS 5977 or later. For systems running Enginuity 5876, the health score is based on the Hardware, Configuration, Capacity, and SG Response time scores. The health score is calculated every five minutes.

**Note:** The Health score values for Hardware, SG Response, and service level compliance are not real time; they are based on values within the last hour.

The **Hardware** section shows the director count for Front End, Back End, and SRDF Directors and the available port count on the system. An alert status is indicated through a colored bell beside the title of the highest level alert in that category. If no alerts are present, then a green tick is displayed.

#### *Replication dashboard view for a specific storage system*

The Replication Dashboard provides storage group summary protection information, summarizing the worst states of various replication technologies and counts of management objects participating in these technologies. For systems running HYPERMAX OS 5977 and higher, summary information for SRDF, SRDF/ Metro, and SnapVX (including zDP snapshots) is displayed. For systems running Enginuity OS 5876, summary information for SRDF and device groups is displayed.

The Replication Dashboard has a SRDF topology view that visually describes the layout of the SRDF connectivity of the selected storage system in Unisphere.

The Replication Dashboard provides a Migrations Environments topology view that visually describes the layout of the migration environments of the selected storage system.

*Storage Group Compliance dashboard view for a specific storage system*

The Storage Group Compliance dashboard displays how well the workload of the storage system is complying with the overall service level. Storage group compliance information displays for storage systems that are registered with the Performance component. The total number of storage groups is listed, along with information about the number of storage groups performing in accordance with service level targets. A list view of the storage groups is also provided and this can be filtered.

*Capacity dashboard view for a specific storage system*

The storage system capacity dashboard enables you to see the amount of capacity your storage system is subscribed for, and the amount of that subscribed capacity that has been allocated. You can also see how efficient the storage system is in using data reduction technologies.

The SRP capacity dashboard reports the capacity and efficiency breakdown of a SRP. For PowerMaxOS 5978 storage systems running 9.1, FBA and CKD devices can be configured in a single SRP. This reduces the cost of storage array ownership for a mixed system and enables the efficient management of drive slot consumption in the array. Where the SRP is of mixed emulation, you can select by emulation to examine breakdown.

*Performance and Capacity dashboard view or a specific storage system*

The performance and capacity dashboard for a specific storage system provides a view of key performance and capacity indicators.

- A **Capacity** panel displays the following:

**Subscribed Capacity**

A graphical representation of the subscribed capacity of the system (used = blue and free = gray) and the percentage used

**Usable Capacity**

A graphical representation of the usable capacity of the system (used = blue and free = gray) and the percentage used

**Subscribed Usable Capacity**

The percentage of subscribed usable capacity

**Overall Efficiency**

The overall efficiency ratio

**Trend**

A panel displaying usable capacity and subscribed capacity in terabytes

- A **Performance** panel displays the following graphs over a four hour, one week, or two-week period:
  - **Host IOs/sec** in terms of read and write operations over time.
  - **Latency** in terms of read and write operations over time.
  - **Throughput** in terms of read and write operations over time.

To the right of each graph, a list of the top five active storage groups for that graph is displayed. Zooming in to a timeframe on a graph automatically updates the top five storage groups lists for that timeframe. Clicking a particular point in time on one graph automatically updates the top five storage group lists for that particular time.

## Understanding the system health score

The **System Health** dashboard provides a single place from which you can quickly determine the health of the system.

The System Health panel displays values for the following high-level health or performance metrics: Configuration, Capacity, System Utilization, Storage Group (SG) response time, and Service Level compliance. It also displays an overall health score that is based on the lowest health score out of the five metrics. These five categories are for storage systems running HYPERMAX OS 5977 or higher. For storage systems running Enginuity OS 5876, the health score is based on four categories: Configuration, Capacity, System Utilization, and Storage Group (SG) response time. The health score is calculated every five minutes. The overall value is always calculated from all metric values. If a health score category is seen as stale or unknown, then the overall health score is not updated. The previously calculated overall health score is displayed but its value is denoted as stale by setting the menu item to grey.

The System Utilization, Capacity, Storage Group response time, and Service Level compliance are based on performance information.

The Configuration health score is calculated every five minutes and is based on the director and port alerts in the system at the time of calculation. Unisphere does not support alert correlation or auto clearing, so you are required to manually delete alerts that have been dealt with or are no longer relevant as these impact on the hardware health score until they are removed from Unisphere.

The Configuration health score is calculated as follows:

- Director out of service- 40 points
- Director Offline - 20 points
- Port Offline - 10 points

The Capacity health score is based on % Effective Used Capacity. For storage systems running HYPERMAX OS 5977 or higher, capacity levels are checked at the Storage Resource Pool (SRP) level. For storage systems running PowerMaxOS 5978 or higher, capacity levels are checked at the SRP level or SRP Emulation level (where a mixed SRP emulation is involved). For storage systems running Enginuity OS 5876, capacity levels are checked at thin pool level.

The capacity health scores are calculated as follows:

- Fatal level - based on what is defined in the **System Threshold and Alerts** dialog. The default fatal threshold is 100% - 30 points.
- Critical level - based on what is defined in the **System Threshold and Alerts** dialog. The default critical threshold is 80% - 20 points.

The System Utilization health score is calculated using the threshold limits of the following categories and metrics:

- FE Director : % busy, queue depth utilization
- FE port: % busy
- BE port: % busy
- BE Director (DA): % busy
- SRDF port: % busy
- SRDF Director: % busy
- DX port: % busy
- External Director: % busy



- IM Director: % busy
- EDS Director: % busy
- Board: Utilization
- Cache Partition: %WP utilization
- Disk: % busy

For each instance and metric for particular category, the threshold info is found. If not set, use the default thresholds. The default thresholds are:

FE Port: % busy - Critical 70, Warning 50

FE Director: % busy - Critical 70, Warning 50; Queue Depth Utilization - Critical 75, Warning 60

BE Port: % busy - Critical 70, Warning 55

BE Director (DA): % busy - Critical 70, Warning 55

SRDF Port: % busy - Critical 70, Warning 50

SRDF Director: % busy - Critical 70, Warning 50

DX Port: - % busy - Critical 70, Warning 55

External Director: % busy- Critical 70, Warning 55

IM Director: % busy - Critical 70, Warning 55

EDS Director: % busy - Critical 70, Warning 55

Board - Utilization - Critical 70, Warning 60

Cache Partition: %WP utilization - Critical 75, Warning 55

Disk: % busy - - Critical 70, Warning 55

The system utilization score is calculated as follows:

- Critical level: - five points

The Storage Group Response health score is based on software category health scores. Certain key metrics are examined against threshold values and if they exceed a certain threshold, the health score is negatively affected.

The storage group response score is based on the following:

- Storage Group: Read Response Time, Write Response Time, Response Time
- Database: Read Response Time, Write Response Time, Response Time

For each instance and metric for particular category, the threshold info is found. If not found, default thresholds are used.

The storage group response score is calculated as follows:

- ■ Read Response Time: Critical - five points
- ■ Write Response Time: Critical - five points
- ■ Response Time: Critical - five points

Storage systems running HYPERMAX OS 5977 or higher: The Service Level Compliance health score is based on Workload Planner (WLP) workload state. A reduction from the health score is performed when storage groups that have a service level that is defined are not meeting the service level requirements.

The Service Level compliance score is calculated as follows:

- Underperforming: - five points

## Managing jobs

### About this task

When you perform certain configuration tasks on the storage system, they are not immediately processed, but are kept in a job list for you to review and submit in batches. One way to identify these tasks is from the dialog boxes. They have a button that is named **Add to Job List**.

Unisphere includes a job list view, from which you can view and manage the job list for a storage system.

## Unisphere for PowerMax alert monitoring recommendations

This section outlines the list of recommended alerts for you to monitor or consider monitoring (depending on your environment) when configuring alert policies using Unisphere for PowerMax.

Note that this relates to storage system running HYPERMAX OS 5977 or higher.

Alert notifications should also be enabled for these alerts.

In addition, notifications should be configured for the default SystemThresholds Alerts set and Notifications set.

It is recommended that you monitor the following:

- Array Component events
- Array Events
- Array - Deferred Service Threshold Alert
- Array - Director Status
- Array - Disk Status
- Array - Environmental Alert
- Array - Hotspare Invoked
- Array - Migration Complete Alert
- Array - Port Link status
- Array - Port status
- Array - RVA Spare Coverage
- Array - SP Alerts
- Array - SRDF Alerts
- Array - SRDF Job Flow Control Change
- Array - SRDF Link Status
- Array - SRDF/A No Cycle Switch Alert
- Array - SRDF/A Session
- Array - SRDF/A Session dropped, transmit idle state timeout
- Array - SRDF/A Session entering transmit idle state
- Array - SRDF/A Session recovered from a transmit idle state
- Array - Target Engenuity Warning

Consider monitoring the following depending on customer environment:

- Array - Device Config Change
- Array - Device Status
- Array - Thin Device Allocation
- Array - Thin Device Usage

## Server alerts

Server alerts are alerts that are generated by Unisphere itself.

Unisphere generates server alerts under the conditions that are listed in the following table:

Checks are run on 10 minute intervals and alerts are raised on 24-hour intervals from the time the server was last started. These time intervals also apply to discover operations. That is, performing a discover operation does not force the delivery of these alerts.

**Note:** Runtime alerts are not storage system-specific. They can be deleted if the user has admin or storage admin rights on at least one storage system. A user with a monitor role cannot delete the server alerts.

Server alert	Number of volumes	Threshold	Alert Details
Total memory on the Unisphere server	0–64,000	12 GB	System memory <# GB> is below the minimum requirement of <# GB>
	64,000–128,000	16 GB	
	128,000–256,000	20 GB	
Free disk space on the Unisphere installed directory	0–64,000	100 GB	Free disk space <# GB> is below the minimum requirement of <# GB>
	64,000–128,000	140 GB	
	128,000–256,000	180 GB	
Number of managed storage systems	Threshold is 20.		Number of managed arrays <#> is over the maximum supported number of #
Number of managed volumes	256,000		Number of managed volumes <#> is over the maximum supported number of <#>. Solutions Enabler may indicate a slightly different number of volumes than indicated in this alert.
Number of gatekeepers	6		Number of gatekeepers <#> on storage system is below the minimum requirement of 6.

## Understanding settings

Systems settings are managed from a central point.

The following categories of settings can be modified:

- Preferences—General and Performance settings
- System and Licenses—License Usage, Solutions Enabler, and System Entitlements settings

- Users and Groups—Authentication, Local Users, User Sessions, and Authorized Users and Groups settings
- System Access Control—Access Control entries, Access groups, and Access Pools settings
- Management—System Attributes, Link and Launch, Secure Remote Services, and CloudIQ settings
- Data Protection—Data Protection settings
- Performance—System Registrations, Dashboard Catalog, Real-Time Traces, Metrics, and Export PV settings
- Unisphere Databases—Performance Databases and System Database settings
- DSA Environment—Database Storage analyzer (DSA) settings
- Alerts—Alerts Policies, Compliance Alert Policies, Performance Thresholds and Alerts, System Thresholds and Alerts, and Notifications settings

Unisphere 9.1 provides the ability of saving specific settings on one array, so that these settings can be applied to other arrays of the same family and PowerMax version. Settings can be cloned, imported, and exported.

## Understanding licenses

Unisphere supports electronic licensing (eLicensing). eLicensing is a license management solution to help you track and comply with software license entitlement.

eLicensing uses embedded locking functions and back-office IT systems and processes. It provides you with better visibility into software assets, easier upgrade, and capacity planning and reduced risk of non-compliance, while still adhering to a strict “do no harm” policy to your operations.

When installing licenses with eLicensing, you obtain license files from customer service, copy them to a Solutions Enabler or a Unisphere host, and load them onto storage systems.

Each license file fully defines the entitlements for a specific system, including its activation type (Individual or Enterprise), the licensed capacity, and the date the license was created. If you want to add a product title or increase the licensed capacity of an entitlement, obtain a new license file from online support and load it onto the storage system.

When managing licenses, Solutions Enabler, Unisphere, z/OS Storage Manager (EzSM), MF SCF native command line, TPF, and IBM i platform console, provide detailed usage reports that enable you to better manage capacity and compliance planning.

There are two types of eLicenses: host-based and array-based. Host-based licenses, as the name implies, are installed on the host. And, array-based licenses are installed on the storage system. For information about the types of licenses and the features they activate, see the *Solutions Enabler Installation Guide*.

Unisphere enables you to add and view array-based licenses, and add, view, and remove host-based licenses.

Unisphere uses array-based eLicensing.

### Note:


For more information about eLicensing, see the *Solutions Enabler Installation Guide*.

## Understanding user authorization

User authorization is a tool for restricting the management operations users can perform on a storage system or with the Database Storage Analyzer application. By default, user authorization is enabled for Unisphere users, regardless of whether it is enabled on the storage system.

When configuring user authorization, an Administrator or SecurityAdmin maps individual users or groups of users to specific roles on storage systems or Database Storage Analyzer, which

determine the operations the users can perform. These user-to-role-to-storage system/Database Storage Analyzer mappings (known as authorization rules) are maintained in the symauth users list file, which is located on either a host or storage system, depending on the storage operating environment.

 **Note:** If there is one or more users listed in the symauth file, users not listed in the file are unable to access or even see storage systems from the Unisphere console.

## Individual and group roles

Users gain access to a storage system or component either directly through a role assignment and/or indirectly through membership in a user group that has a role assignment. If a user has two different role assignments (one as an individual and one as a member of a group), the permissions assigned to the user will be combined. For example, if a user is assigned a Monitor role and a StorageAdmin role through a group, the user is granted Monitor and StorageAdmin rights.

## Roles

The following lists the available roles. Note that you can assign up to four of these roles per authorization rule. For a more detailed look at the permissions that go along with each role, see [Roles and associated permissions](#) on page 15.

- **None**—Provides no permissions.
- **Monitor**—Performs read-only (passive) operations on a storage system excluding the ability to read the audit log or access control definitions.
- **StorageAdmin**—Performs all management (active or control) operations on a storage system and modifies GNS group definitions in addition to all Monitor operations
- **Administrator**—Performs all operations on a storage system, including security operations, in addition to all StorageAdmin and Monitor operations.
- **SecurityAdmin**—Performs security operations on a storage system, in addition to all Monitor operations.
- **Auditor**—Grants the ability to view, but not modify, security settings for a storage system, (including reading the audit log, symacl list and symauth) in addition to all Monitor operations. This is the minimum role required to view the storage system audit log.
- **DSA Admin**—Collects and analyzes database activity with Database Storage Analyzer. A user cannot change their own role so as to remove Administrator or SecurityAdmin privileges from themselves.
- **Local Replication**—Performs local replication operations (SnapVX or legacy Snapshot, Clone, BCV). To create Secure SnapVX snapshots a user needs to have Storage Admin rights at the array level. This role also automatically includes Monitor rights.
- **Remote Replication**—Performs remote replication (SRDF) operations involving devices and pairs. Users can create, operate upon or delete SRDF device pairs but can't create, modify or delete SRDF groups. This role also automatically includes Monitor rights.
- **Device Management**—Grants user rights to perform control and configuration operations on devices. Note that Storage Admin rights are required to create, expand or delete devices. This role also automatically includes Monitor rights.

In addition to these user roles, Unisphere includes an administrative role, the Initial Setup User. This user, defined during installation, is a temporary role that provides administrator-like permissions for the purpose of adding local users and roles to Unisphere.

## User IDs

Users and user groups are mapped to their respective roles by IDs. These IDs consist of a three-part string in the form:

Type:Domain\Name

Where:

- **Type**—Specifies the type of security authority used to authenticate the user or group. Possible types are:
  - **L** — Indicates a user or group authenticated by LDAP. In this case, Domain specifies the domain controller on the LDAP server. For example:  
L:danube.com\Finance  
Indicates that user group Finance logged in through the domain controller danube.com.
  - **C** — Indicates a user or group authenticated by the Unisphere server. For example:  
C:Boston\Legal  
Indicates that user group Legal logged in through Unisphere sever Boston.
  - **H** — Indicates a user or group authenticated by logging in to a local account on a Windows host. In this case, Domain specifies the hostname. For example:  
H:jupiter\mason  
Indicates that user mason logged in on host jupiter.
  - **D** — Indicates a user or group authenticated by a Windows domain. In this case, Domain specifies the domain or realm name. For example:  
D:sales\putman  
Indicates user putman logged in through a Windows domain sales.
- **Name**—specifies the username relative to that authority. It cannot be longer than 32 characters and spaces are allowed if delimited with quotes. Usernames can be for individual users or user groups.

Within role definitions, IDs can be either fully qualified (as shown above), partially qualified, or unqualified. When the Domain portion of the ID string is an asterisk (\*), the asterisk is treated as a wildcard, meaning any host or domain.

When configuring group access, the Domain portion of the ID must be fully qualified.

For example:

- D:ENG\jones—Fully qualified path with a domain and username (for individual domain users).
- D:ENG.xyz.com\ExampleGroup—Fully qualified domain name and group name (for domain groups).
- D:\*jones—Partially qualified that matches username jones with any domain.
- H:HOST\jones—Fully qualified path with a hostname and username.
- H:\*jones—Partially qualified that matches username jones within any host.
- jones—Unqualified username that matches any jones in any domain on any host.

In the event that a user is matched by more than one mapping, the user authorization mechanism uses the more specific mapping. If an exact match (for example, D:sales\putman) is found, that is used; if a partial match (for example, D:\*putman) is found, that is used; if an unqualified match (for example, putman) is found, that is used; otherwise, the user is assigned a role of None.

## Roles and associated permissions

The following tables detail the permissions that go along with each role in Unisphere.

**Note:** The Unisphere Initial Setup User has all permissions on a storage system until an Administrator or SecurityAdmin is added to the storage system.

The roles and the acronyms used for them in these tables are:

- Administrator (AD)
- StorageAdmin (SA)
- Monitor (MO)
- SecurityAdmin (SecA)
- Auditor (AUD)
- None
- PerfMonitor (PM)
- Database Storage Analyzer Admin (DSA)
- Local Replication
- Remote Replication
- Device Management

**Table 1** User roles and associated permissions

Permissions	AD	SA	MO	SecA	AUD	None	PM	DSA
Create/delete user accounts	Yes	No	No	Yes	No	No	No	No
Reset user password	Yes	No	No	Yes	No	No	No	No
Create roles	Yes	Yes	No	Yes (self-excluded)	No	No	No	No
Change own password	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Manage storage systems	Yes	Yes	No	No	No	No	No	No
Discover storage systems	Yes	No	No	Yes	No	No	No	No
Add/show license keys	Yes	Yes	No	No	No	No	No	No
Set alerts and Optimizer monitoring options	Yes	Yes	No	No	No	No	No	No
Release storage system locks	Yes	Yes	No	No	No	No	No	No

**Table 1** User roles and associated permissions (continued)

Permissions	AD	SA	MO	SecA	AUD	None	PM	DSA
Set Access Controls	Yes	Yes	No	No	No	No	No	No
Set replication and reservation preferences	Yes	Yes	No	No	No	No	No	No
View and export the storage system audit log	Yes	No	No	Yes	Yes	No	No	No
Access performance data	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Start data traces	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Set performance thresholds/alerts	Yes	Yes	No	No	No	No	Yes	No
Create and manage performance dashboards	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Collect and analyze database activity with DSA	No	No	No	No	No	No	No	Yes

**Table 2** Permissions for Local Replication, Remote Replication and Device Management roles

Permissions			
	Local Replication	Remote Replication	Device Management
Create/delete user accounts	No	No	No
Reset user password	No	No	No
Create roles	No	No	No
Change own password	Yes	Yes	Yes
Manage storage systems	No	No	No
Discover storage systems	No	No	No
Add/show license keys	No	No	No
Set alerts and Optimizer monitoring options	No	No	No




**Table 2** Permissions for Local Replication, Remote Replication and Device Management roles (continued)

Permissions	Local Replication	Remote Replication	Device Management
	Release storage system locks	No	No
Set Access Controls	No	No	No
Set replication and reservation preferences	No	No	No
View the storage system audit log	No	No	No
Access performance data	Yes	Yes	Yes
Start data traces	Yes	Yes	Yes
Set performance thresholds/alerts	No	No	No
Create and manage performance dashboards	Yes	Yes	Yes
Collect and analyze database activity with Database Storage Analyzer	No	No	No
Perform control and configuration operations on devices	No	No	Yes
Create, expand or delete devices	No	No	No
Perform local replication operations (SnapVX, legacy Snapshot, Clone, BCV)	Yes	No	No
Create Secure SnapVX snapshots	No	No	No
Create, operate upon or delete SRDF device pairs	No	Yes	No
Create, modify or delete SRDF groups	No	No	No

## RBAC roles for TimeFinder SnapVX local and remote replication actions

A user needs to be assigned the necessary roles to perform TimeFinder SnapVX local and remote replication actions.

The following table details the roles that are needed to perform TimeFinder SnapVX local and remote replication actions:

 **Note:** Unisphere for PowerMax does not support RBAC device group management.

	Local Replication	Remote Replication	Device Manager
Protection Wizard - Create SnapVx Snapshot	Yes <sup>(a)</sup>		
Create Snapshot	Yes <sup>(a)</sup>		
Edit Snapshot	Yes		
Link Snapshot	Yes <sup>(b) (c)</sup>		Yes <sup>(d)</sup>
Relink Snapshot	Yes <sup>(b) (c)</sup>		Yes <sup>(d)</sup>
Restore Snapshot	Yes <sup>(b)</sup>		Yes <sup>(b)</sup>
Set Time To Live	Yes		
Set Mode	Yes <sup>(b)</sup>		Yes <sup>(d)</sup>
Terminate Snapshot	Yes		
Unlink Snapshot	Yes <sup>(b)</sup>		Yes <sup>(d)</sup>

(a) - Set Secure is blocked for users who only have Local\_REP rights.

(b) - The user must have the specified rights on the source volumes.


(c) - The user may only choose existing storage groups to link to. Creating a storage group requires Storage Admin rights.

(d) - The user must have the specified rights on the link volumes.

## RBAC roles for SRDF local and remote replication actions

A user needs to be assigned the necessary roles to perform SRDF local and remote replication actions.

The following table details the roles that are needed to perform SRDF local and remote replication actions:

 **Note:** Unisphere for PowerMax does not support RBAC device group management.

	Local Replication	Remote Replication	Device Manager
SRDF Delete		Yes	
SRDF Establish		Yes	
SRDF Failback		Yes	

	Local Replication	Remote Replication	Device Manager
SRDF Failover		Yes	
SRDF Invalidate		Yes	
SRDF Move		Yes	
SRDF Not Ready		Yes	
SRDF R1 Update		Yes	
SRDF Ready		Yes	
SRDF Refresh		Yes	
SRDF Restore		Yes	
SRDF Resume		Yes	
SRDF RW Disable R2		Yes	
SRDF RW Enable		Yes	
SRDF Set Bias		Yes	
SRDF Set Consistency		Yes	
SRDF Set Mode		Yes	
SRDF Set SRDFA		Yes	
SRDF Split		Yes	
SRDF Suspend		Yes	
SRDF Swap		Yes	
SRDF Write Disable		Yes	

## Understanding access controls for volumes

Access controls can be set on specific volumes within a storage system and those volumes can be assigned to a specific host.

Administrators, StorageAdmins, and SecurityAdmins can set access controls on specific volumes within a storage system and assign those volumes to a specific host. When set, only that host can see the volumes, and perform the granted operations. Other hosts that are connected to that storage system do not see those volumes. This behavior eliminates the possibility of one host inadvertently performing operations on volumes that belong to someone else.

 **Note:** Refer to the *Solutions Enabler Array Management CLI Product Guide* for more information about Access Controls.

## Understanding Storage Management

Overview of storage management.

Storage Management covers the following areas:

- Storage Group management - Storage groups are a collection of devices that are stored on the array, and an application, a server, or a collection of servers use them. Storage groups are used

to present storage to hosts in masking/mapping, Virtual LUN Technology, FAST, and various base operations.

- Service Levels management - A service level is the response time target for a storage group. The service level sets the storage array with the required response time target for a storage group. It automatically monitors and adapts to the workload needed maintain the response time target. The service level includes an optional workload type so it can be optimized to meet performance levels.
- Templates management - Using the configuration and performance characteristics of an existing storage group as a starting point, you can create templates that will prepopulate fields in the provisioning wizard and create a more realistic performance reservation in your future provisioning requests.
- Storage Resource Pools management - Fully Automated Storage Tiering (FAST) provides automated management of storage array disk resources to achieve expected service levels. FAST automatically configures disk groups to form a Storage Resource Pool (SRP) by creating thin pools according to each individual disk technology, capacity, and RAID type.
- Volumes management - A storage volume is an identifiable unit of data storage. Storage groups are sets of volumes.
- External Storage management - Fully Automated Storage Tiering - External (FAST.X) attaches external storage to storage systems directs workload movement to these external arrays while having access to the array features such as local replication, remote replication, storage tiering, data management, and data migration. Also, it simplifies multi-vendor or Dell EMC storage array management.
- VVol management - VMware VVols enable data replication, snapshots, and encryption to be controlled at the VMDK level instead of the LUN level, where these data services are performed on a per VM (application level) basis from the storage array.
- FAST Policies management - A FAST policy consists of one to three DP tiers, or one to four VP tiers, but not a combination of both DP and VP tiers. Policies define a limit for each tier in the policy. This limit determines the amount of data from a storage group that is associated with the policy that can reside on the tier.
- Tiers management - FAST automatically moves active data to high-performance storage tiers and inactive data to low-cost, high-capacity storage tiers.
- Thin Pools management - Storage systems are preconfigured at the factory with virtually provisioned devices. Thin Provisioning helps reduce cost, improve capacity utilization, and simplify storage management. Thin Provisioning presents a large amount of capacity to a host and then consumes space only as needed from a shared pool. Thin Provisioning ensures that thin pools can expand in small increments while protecting performance, and performs nondisruptive shrinking of thin pools to help reuse space and improve capacity utilization.
- Disk Groups management - A disk group is a collection of hard drives within the storage array that share the same performance characteristics.
- VLUN Migration management - Virtual LUN Migration (VLUN Migration) enables transparent, nondisruptive data mobility for disk group provisioned storage system volumes, virtually provisioned storage system volumes, between storage tiers and between RAID protection schemes. Virtual LUN can be used to populate newly added drives or move volumes between high performance and high capacity drives, resulting in the delivery of tiered storage capabilities within a single storage system. Migrations are performed while providing constant data availability and protection.

## Understanding storage provisioning

Service level provisioning simplifies storage management by automating many of the tasks that are associated with provisioning storage.

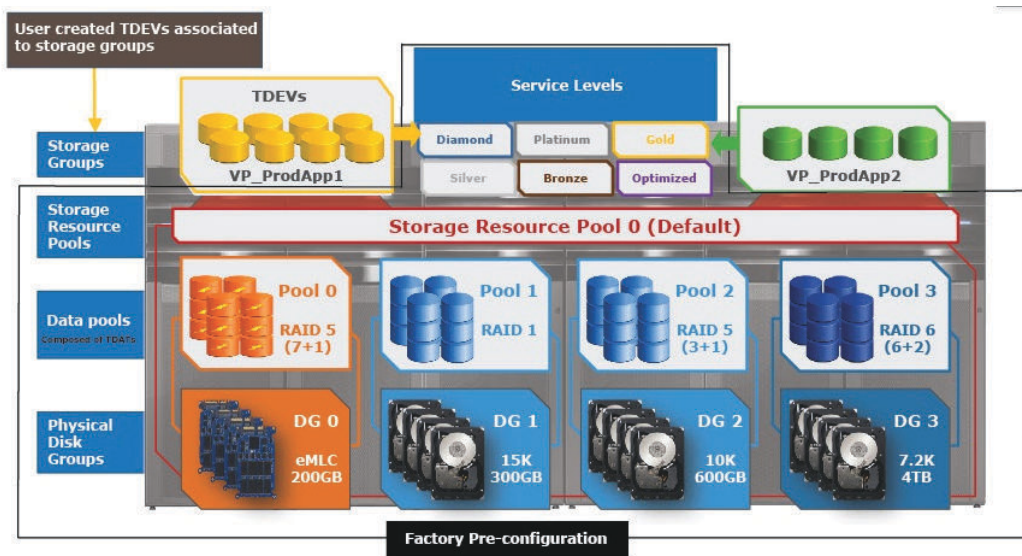
With the release of HYPERMAX OS 5977, Unisphere introduces support for service level provisioning. Service level provisioning simplifies storage management by automating many of the tasks that are associated with provisioning storage.

Service level provisioning eliminates the need for storage administrators to manually assign physical resources to their applications. Instead, storage administrators specify the storage performance and capacity that is required for the application and let the system provision the workload appropriately.

By default, storage systems running HYPERMAX OS 5977 or higher are preconfigured with a single Storage Resource Pool (SRP). The SRP contains all the hard disks on the system that is organized into disk groups by technology, capacity, rotational speed, and RAID protection type. Storage administrators can view all SRPs configured on the system, and the demand that storage groups are placing on them.

Storage systems are also preconfigured with several service levels and workloads. Storage administrators use the service levels and workloads to specify the performance objectives for the application they are provisioning.

When provisioning storage for an application, storage administrators assign the appropriate SRP, service level, and workload to the storage group containing the LUNs associated with the application.



Unisphere provides the following methods for provisioning storage:

**Recommended:** This method relies on wizards to step you through the provisioning process. It is best suited for novice and advanced users who do not require a high level of customization. Customization is the ability to create their own volumes, storage groups, and so on.

**Advanced:** This method, as its name implies, is for advanced users who want the ability to control every aspect of the provisioning process.

This section provides the high-level steps for each method, with links to the relevant help topics for more detail.

Regardless of the method you choose, once you have completed the process, a masking view has been created. In the masking view, the volumes in the storage group are masked to the host initiators and mapped to the ports in the port group.

Before you begin:

The storage system has been configured.

To provision storage for storage systems running HYPERMAX OS 5977 or higher:

<ol style="list-style-type: none"> <li>1. Use the <b>Create Host</b> dialog box to group host initiators (HBAs).</li> <li>2. Use the <b>Provision Storage</b> wizard, which steps you through the process of creating the storage group, port group, and masking view.</li> </ol>	<ol style="list-style-type: none"> <li>1. Use the <b>Create Host</b> dialog box to group host initiators (HBAs).</li> <li>2. Create one or more volumes on the storage system.</li> <li>3. Use the <b>Create Storage Group</b> dialog box to add the created volumes to a storage group, and associate the storage group with a storage resource pool, a service level, and a workload.</li> <li>4. Group Fibre Channel and/or iSCSI front-end directors.</li> <li>5. Associate the host, storage group, and port group into a masking view.</li> </ol>
---	---


Unisphere provides the following methods for provisioning storage on storage systems running Enginuity OS 5876:

**Recommended:** This method relies on wizards to step you through the provisioning process. It is best suited for novice and advanced users who do not require a high level of customization. Customization is the ability to create their own volumes, storage groups, and so on.

**Advanced:** This method, as its name implies, is for advanced users who want the ability to control every aspect of the provisioning process.

This section provides the high-level steps for each method, with links to the relevant help topics for more detail.

Regardless of the method you choose, once you have completed the process, a masking view has been created. In the masking view, the volumes in the storage group are masked to the host initiators and mapped to the ports in the port group.



Before you begin:

The storage system has been configured.

To provision storage for storage systems running Enginuity OS 5876:

<ol style="list-style-type: none"> <li>1. Use the Create Host dialog box to group host initiators (HBAs).</li> <li>2. Use the Provision Storage wizard, which steps you through the process of creating the storage group, port group, and masking view. The wizard optionally associates the storage group with a FAST policy.</li> </ol>	<ol style="list-style-type: none"> <li>1. Use the Create Host dialog box to group host initiators (HBAs).</li> <li>2. Create one or more volumes on the storage system.</li> <li>3. Use the Create Storage Group wizard to create a storage group. If you want to add the volumes you created in step 2, be sure to set the Storage Group Type to Empty, and then complete adding volumes to storage groups.</li> <li>4. Group Fibre Channel and/or iSCSI front-end directors.</li> <li>5. Associate the host, storage group, and port group into a masking view.</li> <li>6. Associate the storage group with a FAST policy.</li> </ol> <p>Optional: Associate the storage that you created in step 3 with an existing FAST policy and assign a priority value for the association.</p>
--	--

## Understanding storage groups

Storage groups are a collection of devices that are stored on the array, and an application, a server, or a collection of servers use them. Storage groups are used to present storage to hosts in masking/mapping, Virtual LUN Technology, FAST, and various base operations.

For storage groups on storage systems running HYPERMAX OS 5977 or higher:

- The maximum number of storage groups that are allowed on a storage system running HYPERMAX OS 5977 is 16,384.
- For HYPERMAX OS 5977 or higher, the maximum number of child storage groups that are allowed in a cascaded configuration is 64.
- A storage group can contain up to 4,096 volumes.
- A volume can belong to multiple storage groups when only one of the groups is under FAST control.
- You cannot create a storage group containing CKD volumes and FBA volumes.

For storage groups on storage systems running Enginuity OS 5876:

- The maximum number of storage groups that are allowed on a storage system running Enginuity OS 5876 is 8,192.
- For Enginuity 5876 or higher, the maximum number of child storage groups that are allowed in a cascaded configuration is 32.

## Understanding compression

Compression allows users to compress user data on storage groups and storage resources. The feature is enabled by default and can be turned on and off at storage group and storage resource level.

If a storage group is cascaded, enabling compression at this level enables compression for each of the child storage groups. The user has the option to disable compression on one or more of the child storage groups if desired.


To turn the feature off on a particular storage group or storage resource, uncheck the **Compression** check box in the in the **Create Storage Group**, **Modify Storage Group** or **Add Storage Resource To Storage Container** dialogs or when using the **Provision Storage** or **Create Storage Container** wizards.

The following are the prerequisites for using compression:

- Compression is only allowed on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher.
- Compression is allowed for FBA devices only.
- The user must have at least StorageAdmin rights.
- The storage group needs to be FAST managed.
- The associated SRP cannot be comprised, either fully or partially, of external storage.

### Reporting

Users are able to see the current compression ratio on the device, the storage group and the SRP. Efficiency ratios are reported in units of 1/10th:1.

 **Note:** External storage is not included in efficiency reports. For mixed SRPs with internal and external storage only the internal storage is used in the efficiency ratio calculations.

## Understanding service levels

A service level is the response time target for the storage group. The service level enables you set the storage array with the desired response time target for the storage group.

It automatically monitors and adapts to the workload to maintain (or meet) the response time target. The service level includes an optional workload type. The optional workload type can be used to further tune expectations for the workload storage group to provide enough flash to meet your performance objective.

## Suitability Check restrictions

The suitability check determines whether the storage system can handle the updated service level.

The Suitability Check option is only available when:

- The storage system is running HYPERMAX OS 5977 or higher.
- The storage system is registered with the performance data processing option for statistics.
- The workloads have been processed.
- All the storage groups that are involved have a service level and SRP set.
- The target SRP does not contain only external disk groups (like XtreamIO).



- The storage system is local.
- The storage group is not in a masking view (only for the local provisioning wizard).
- An issue with one of the selected ports arises when provisioning storage and a valid Front-End Suitability score cannot be derived. Examples of issues are: a virtual port is selected, an offline port is selected, and a selected port has no negotiated speed. When an issue arises, 200.0% (not a real suitability score) is displayed. Excluding data has no impact on the 200% displayed value.

## Understanding storage templates

### About this task

Storage templates are a reusable set of storage requirements that simplify storage management for virtual data centers by eliminating many of the repetitive tasks required to create and make storage available to hosts/applications. With this feature, Administrators and Storage Administrators create templates for their common provisioning tasks and then invoke them later when performing such things as:

- Creating or provisioning storage groups.

The templates created on a particular Unisphere server can be used across all the arrays on that particular server.

Storage templates require storage system running HYPERMAX OS 5977 or greater and storage groups.

A provisioning template contains configuration information and a performance reservation.

A Workload Plan/Performance Reservation is the I/O profile (IOPS/MBPS, Skew Mixture) for a particular SL-WL type combination. By default, the reservation is used for suitability checks and for comparison to current load running. The reservation expires after 14 days.

## Understanding Storage Resource Pools

A Storage Resource Pool is a collection of data pools that provide FAST a domain for capacity and performance management.

By default, a single default Storage Resource Pool is preconfigured at the factory. More Storage Resource Pools can be created with a service engagement. FAST performs all its data movements within the boundaries of the Storage Resource Pool.

## Understanding volumes

A storage volume is an identifiable unit of data storage. Storage groups are sets of volumes.

The Volumes view on the Unisphere user interface provides you with a single place from which to view and manage all the volume types on the system.

## Understanding Federated Tiered Storage

Federated Tiered Storage (FTS) enables you to attach external storage to a storage system.


Attaching external storage enables you to use physical disk space on existing storage systems. You also gain access to features such as local replication, remote replication, storage tiering, data management, and data migration.

For additional information about FTS, see the following documents:

- *Federated Tiered Storage (FTS) Technical Notes*
- *Solutions Enabler Array Management CLI Product Guide*
- *Solutions Enabler TimeFinder Family CLI User Guide*

## Understanding FAST

Fully Automated Storage Tiering (FAST) automates management of storage system disk resources on behalf of thin volumes.

 **Note:** This section describes FAST operations for storage systems running HYPERMAX OS 5977 or higher.

FAST automatically configures disk groups to form a Storage Resource Pool by creating thin pools according to each individual disk technology, capacity, and RAID type.

FAST technology moves the most active parts of your workloads (hot data) to high-performance flash disks and the least-frequently accessed storage (cold data) to lower-cost drives, using the best performance and cost characteristics of each different drive type. FAST delivers higher performance using fewer drives to help reduce acquisition, power, cooling, and footprint costs. FAST can factor in the RAID protections to ensure write heavy workloads go to RAID 1 and read heavy workloads go to RAID 6. This process is entirely automated and requires no user intervention.

FAST also delivers variable performance levels through service levels. Thin volumes can be added to storage groups and the storage group can be associated with a specific service level to set performance expectations.

FAST monitors the performance of the storage group relative to the service level and automatically provisions the appropriate disk resources to maintain a consistent performance level.

## Understanding Workload Planner

Workload Planner is a FAST component that is used to display performance metrics for applications. It also models the impact of migrating the workload from one storage system to another.

Workload Planner is supported on storage systems running Enginuity OS 5876 or HYPERMAX OS 5977.

For storage groups to be eligible for workload planning, they must meet the following criteria:

- On a locally attached storage system registered for performance
- Belong to only one masking view.
- Under FAST control:
  - For storage systems running HYPERMAX OS 5977, they must be associated with a service level.
  - For storage systems running Enginuity OS 5876, they must be associated with a FAST policy.
- Contain only FBA volumes.

Also, the Unisphere server must be on an open systems host.

## Understanding time windows

Time windows are used by FAST, FAST VP, and Optimizer to specify when data can be collected for performance analysis and when moves/swaps can execute.

There are two types of time windows:

- **Performance time windows**—Specify when performance samples can be taken for analysis.
- **Move time windows**—Specify when moves/swaps are allowed to start or not start.

In addition, performance and move time windows can be further defined as open or closed:

- **Open**—When creating performance time windows, this specifies that the data collected in the time window should be included in the analysis. When creating move time windows, this specifies that the moves can start within the time window. This type of time window is also referred to as inclusive.
- **Closed**—When creating performance time windows, this specifies that the data collected in the time window should be excluded from analysis. When creating move time windows, this specifies that the moves cannot start within the time window. This type of time window is also referred to as exclusive.

## Understanding FAST.X

FAST.X enables the integration of storage systems running HYPERMAX OS 5977 or higher and heterogeneous arrays.

FAST.X enables LUNs on external storage to be used as raw capacity. Data services such as SRDF, TimeFinder, and Open Replicator are supported on the external device.

For additional information, see the following documents:

- *Solutions Enabler Array Management CLI Guide*
- *Solutions Enabler TimeFinder CLI User Guide*

## Virtualize external LUNs

When you attach external storage to a storage system, FAST.X virtualizes the SCSI logical units of an external storage system as disks called eDisks.

eDisks have two modes of operation:

### Encapsulation

Allows you to preserve existing data on external storage systems and access it through storage volumes. These volumes are called encapsulated volumes.

### External Provisioning

Allows you to use external storage as raw capacity for new storage volumes. These volumes are called externally provisioned volumes. Existing data on the external volumes is deleted when they are externally provisioned.

The following restrictions apply to eDisks:

- Can only be unprotected volumes. The RAID protection scheme of eDisks depends on the external storage system.
- Cannot be AS400, CKD, or gatekeeper volumes.
- Cannot be used as VAULT, SFS, or ACLX volumes.

### Encapsulation

Encapsulation has two modes of operation:

#### Encapsulation for disk group provisioning (DP encapsulation)

The eDisk is encapsulated and exported from the storage system as disk group provisioned volumes.


### Encapsulation for virtual provisioning (VP encapsulation)

The eDisk is encapsulated and exported from the storage system as thin volumes.

In either case, Enginuity automatically creates the necessary volumes. If the eDisk is larger than the maximum volume capacity or the configured minimum auto meta capacity, Enginuity creates multiple volumes to account for the full capacity of the eDisk. These volumes are concatenated into a single concatenated meta volume to enable access to the complete volume of data available from the eDisk.

### External provisioning

After you virtualize an eDisk for external provisioning, you can create volumes from the external disk group and present the storage to users. You can also use this storage to create a new FAST VP tier.

 **Note:** If you use external provisioning, any data that is on the external volume is deleted.

### Geometry of encapsulated volumes

Storage volumes are built based on the storage system cylinder size (fifteen 64 K tracks), so the capacity of storage volumes does not always match the raw capacity of the eDisk. If the capacity does not match, Enginuity sets a custom geometry on the encapsulated volume. For created meta volumes, Enginuity defines the geometry on the meta head, and only the last member can have a capacity that spans beyond the raw capacity of the eDisk.

Encapsulated volumes that have a cylinder size larger than the reported user-defined geometry are considered geometry limited. For more details and a list of restrictions that apply to geometry-limited volumes, see the *Solutions Enabler Array Controls CLI Guide*.

## Understanding tiers

FAST automatically moves active data to high-performance storage tiers and inactive data to low-cost, high-capacity storage tiers.

The following rules apply to tier creation:

- This feature requires Enginuity OS 5876.
- The maximum number of tiers that can be defined on a storage system is 256.
- When a disk group or thin pool is specified, its technology type must match the tier technology.
- Disk groups can only be specified when the tier include type is static.
- A standard tier cannot be created if it :
  - Leads to static and dynamic tier definitions in the same technology.
  - Partially overlaps with an existing tier. Two tiers partially overlap when they share only a subset of disk groups. For example, Tier A partially overlaps with Tier B when Tier A contains disk groups 1 and 2, and Tier B contains only disk group 2.

## Understanding thin pools

Storage systems are preconfigured at the factory with virtually provisioned devices. Thin Provisioning helps reduce cost, improve capacity utilization, and simplify storage management. Thin Provisioning presents a large amount of capacity to a host and then consumes space only as needed from a shared pool. Thin Provisioning ensures that thin pools can expand in small increments while protecting performance, and performs nondisruptive shrinking of thin pools to help reuse space and improve capacity utilization.

Unisphere works on a best effort basis when creating thin pools, meaning that it attempts to satisfy as much as possible of the requested pool from existing DATA volumes, and then creates the volumes necessary to meet any shortfall.

Before you begin:

Thin pools contain DATA volumes of the same emulation and the same configuration.

When creating thin pools, Unisphere attempts to instill best practices in the creation process by updating the default protection level according to the selected disk technology:

Technology	Default protection level
EFD	RAID5(3+1)
FC	2-Way Mirror
SATA	RAID6(6+2)


## Understanding disk groups

A disk group is a collection of hard drives within the storage array that share the same performance characteristics.

Disk groups can be viewed and managed from the Unisphere user interface.

## Understanding Virtual LUN Migration

Virtual LUN Migration (VLUN Migration) enables transparent, nondisruptive data mobility for both disk group provisioned and virtually provisioned storage system volumes between storage tiers and between RAID protection schemes. Virtual LUN can be used to populate newly added drives or move volumes between high performance and high capacity drives, thereby delivering tiered storage capabilities within a single storage system. Migrations are performed while providing constant data availability and protection.

 **Note:** Virtual LUN migration requires Engenuity 5876.

Virtual LUN Migration performs tiered storage migration by moving data from one RAID group to another, or from one thin pool to another. It is also fully interoperable with all other storage system replication technologies such as SRDF, TimeFinder/Clone, TimeFinder/Snap, and Open Replicator.

RAID Virtual Architecture allows, for the purposes of migration, two distinct RAID groups, of different types or on different storage tiers, to be associated with a logical volume. In this way, Virtual LUN allows for the migration of data from one protection scheme to another, for example RAID 1 to RAID 5, without interruption to the host or application accessing data on the storage system volume.

Virtual LUN Migration can be used to migrate regular storage system volumes and metavolumes of any emulation — FBA, CKD, and IBM i series. Migrations can be performed between all drive types including high-performance enterprise Flash drives, Fibre Channel drives, and large capacity SATA drives.

Migration sessions can be volume migrations to configured and unconfigured space, or migration of thin volumes to another thin pool.

## Understanding vVols

VMware VVols enables data replication, snapshots, and encryption to be controlled at the VMDK level instead of the LUN level, where these data services are performed on a per VM (application level) basis from the storage array.

The **VVol Dashboard** provides a single place to monitor and manage VVols.

The storage system must be running HYPERMAX OS 5977 or higher.

## Understanding Host Management

Host Management covers the following areas:

- Management of host and host groups
- Management of masking views - A masking view is a container of a storage group, a port group, and an initiator group, and makes the storage group visible to the host. Devices are masked and mapped automatically. The groups must contain some devices entries.
- Management of port groups - Port groups contain director and port identification and belong to a masking view. Ports can be added to and removed from the port group. Port groups that are no longer associated with a masking view can be deleted.
- Management of initiators and initiator groups - An initiator group is a container of one or more host initiators (Fibre or iSCSI). Each initiator group can contain up to 64 initiator addresses or 64 child IG names. Initiator groups cannot contain a mixture of host initiators and child IG names
- Monitor of Xtrem SW Cache (host) cache adapters.
- Management of PowerPath hosts
- Management of mainframe configured splits, CU images, and CKD volumes

## Understanding hosts

Storage hosts are systems that use storage system LUN resources. A logical unit number (LUN) is an identifier that is used for labeling and designating subsystems of physical or virtual storage.

- The maximum number of initiators that are allowed in a host depends on the storage operating environment:
  - For Enigunity OS 5876, the maximum number of initiators that is allowed is 32.
  - For HYPERMAX OS 5977 or higher, the maximum number of initiators that is allowed is 64.

## Understanding masking views

A masking view is a container of a storage group, a port group, and an initiator group, and makes the storage group visible to the host.

Masking views are manageable from the Unisphere user interface. Devices are masked and mapped automatically. The groups must contain some devices entries.

## Understanding port groups

Port groups contain director and port identification and belong to a masking view. Ports can be added to and removed from the port group. Port groups that are no longer associated with a masking view can be deleted.

Note the following recommendations:

Port groups should contain four or more ports.

Each port in a port group should be on a different director.

A port can belong to more than one port group. However, for storage systems running HYPERMAX OS 5977 or higher, you cannot mix different types of ports (physical FC ports, virtual ports, and iSCSI virtual ports) within a single port group.

## Understanding initiators

An initiator group is a container of one or more host initiators (Fibre or iSCSI).

Each initiator group can contain up to 64 initiator addresses or 64 child IG names. Initiator groups cannot contain a mixture of host initiators and child IG names.

## Understanding PowerPath hosts

PowerPath is host-based software that provides automated data path management and load-balancing capabilities for heterogeneous server, network, and storage deployed in physical and virtual environments.

The following are the minimum requirements to perform this task:

- A storage system running PowerMax 5978 or higher
- PowerPath 6.3

## Understanding mainframe management

Service level provisioning for mainframe simplifies storage system management by automating many of the tasks that are associated with provisioning storage.

The mainframe dashboard provides you with a single place to monitor and manage configured splits, CU images, and CKD volumes.

The mainframe dashboard is organized into the following panels:

- **CKD Compliance** - Displays how well CKD storage groups are complying with their respective service level policies, if applicable.
- **CKD Storage Groups** - Displays the mainframe storage groups on the array. Double-click a storage group to see more details and information about its compliance and volumes.
- **Actions** - Allow the user to provisioning storage and create CKD volumes.
- **Summary** - Displays the mainframe summary information in terms of splits, CU images, and CKD volumes.

With the release of HYPERMAX OS 5977 Q1 2016, Unisphere introduces support for service level provisioning for mainframe. Service level provisioning simplifies storage system management by automating many of the tasks that are associated with provisioning storage.

Service level provisioning eliminates the need for storage administrators to manually assign physical resources to their applications. Instead, storage administrators specify the service level

and capacity that is required for the application and the system provisions the storage group appropriately.

You can provision CKD storage to a mainframe host using the Provision Storage wizard.

The storage system must be running HYPERMAX OS 5977 Q1 2016, or higher, and have at least one FICON director configured.

You can map CKD devices to front-end EA/EF directors. Addressing on EA and EF directors is divided into Logical Control Unit images, also known as CU images. Each CU image has its own unique SSID and contains a maximum of 256 devices (numbered 0x00 through 0xFF). When mapped to an EA or EF port, a group of devices becomes part of a CU image.

With the release of HYPERMAX OS 5977 Q2 2017, Unisphere introduces support for All Flash Mixed FBA/CKD arrays.

**Note:** This feature is only available for All Flash 450F/850F/950F arrays that are:

- Purchased as a mixed All Flash system
- Installed at HYPERMAX OS 5977 Q2 2017 or later
- Configured with two Storage Resource Pools - one FBA Storage Resource Pool and one CKD Storage Resource Pool

You can provision FBA/CKD storage to a mainframe host using the **Provision Storage** wizard.

**Note:**

1. A CKD SG can only provision from a CKD SRP.
2. A FBA SG can only provision from a FBA SRP.
3. FBA volumes cannot reside in a CKD SRP.
4. CKD volumes cannot reside in a FBA SRP.
5. Compression is only for FBA volumes.

You can map FBA devices to front-end EA/EF directors. Addressing on EA and EF directors is divided into Logical Control Unit images (CU images). Each CU image has its own unique SSID and contains a maximum of 256 devices (numbered 0x000 through 0xFF). When mapped to an EA or EF port, a group of devices becomes part of a CU image.

## Understanding Data Protection Management

Data Protection management ensures that data is protected and remains available.

Data Protection Management covers the following areas:

- Management of SRDF protected storage groups
- Management of device groups - A device group is a user-defined group consisting of devices that belong to a locally attached array. Control operations can be performed on the group as a whole, or on the individual device pairs in the group. By default, a device can belong to more than one device group.
- Management of SRDF groups - SRDF groups provide a collective data transfer path linking volumes of two separate storage systems. These communication and transfer paths are used to synchronize data between the R1 and R2 volume pairs that are associated with the RDF group. At least one physical connection must exist between the two storage systems within the fabric topology. See *Dell EMC SRDF Introduction* for an overview of SRDF.
- Management of Non-disruptive migration (NDM) management - NDM enables you to migrate storage group (application) data in a non-disruptive manner with no downtime from NDM capable source arrays to NDM capable target arrays.



- Management of SRDF/A DSE Pools
- Management of TimeFinder Snap pool
- Management of Open Replicator - Open Replicator is a software tool that is used to migrate data from third-party arrays to PowerMax.
- Management of RecoverPoint systems
- Management of Virtual Witness - The Witness feature supports a third party that the two storage systems consult when they lose connectivity with each other, that is, their SRDF links go out of service. When SRDF links go out of service, the Witness helps to determine, for each SRDF/Metro Session, which of the storage systems should remain active (volumes continue to be read and write to hosts) and which goes inactive (volumes not accessible).

## Managing remote replication sessions

The SRDF dashboard provides a single place to monitor and manage SRDF sessions on a storage system. This includes device groups types R1, R2, and R21.

See *Dell EMC SRDF Introduction* for an overview of SRDF.

Unisphere provides the ability to monitor and manage the SRDF replication on storage groups directly without the need to map to a device group.

Unisphere provides the ability to monitor and manage SRDF/Metro from the SRDF dashboard. SRDF/Metro delivers active-active high availability for non-stop data access and workload mobility – within a data center and across metro distance. It provides array clustering for storage systems running HYPERMAX OS 5977 or higher enabling even more resiliency, agility, and data mobility. SRDF/Metro enables hosts and host clusters to directly access a LUN or storage group on the primary SRDF array and secondary SRDF array (sites A and B). This level of flexibility delivers the highest availability and best agility for rapidly changing business environments.

In a SRDF/Metro configuration, SRDF/Metro utilizes the SRDF link between the two sides of the SRDF device pair to ensure consistency of the data on the two sides. If the SRDF device pair becomes Not Ready (NR) on the SRDF link, SRDF/Metro must respond by choosing one side of the SRDF device pair to remain accessible to the hosts, while making the other side of the SRDF device pair inaccessible. There are two options which enable this, Bias and Witness.

The first option, Bias, is a function of the two storage systems running HYPERMAX OS 5977 taking part in the SRDF/Metro and is a required and integral component of the configuration. The second option, Witness, is an optional component of SRDF/Metro which allows a third storage system running Engenuity 5876 or HYPERMAX OS 5977 system to act as an external arbitrator to avoid an inconsistent result in cases where the bias functionality alone may not result in continued host availability of a surviving non-biased array.

## Understanding non-disruptive migration

Non-Disruptive Migration (NDM) provides a method for migrating data from a source array to a target array without application host downtime. Minimally disruptive migration enables migrations on the same supported platforms as NDM it but requires a short application outage.

Click [here](#) and search for NDM to watch a video that provides an overview of NDM. Other NDM information is available in the *Dell EMC Solutions Enabler Array Controls and Management CLI User Guide* and the *Non-Disruptive Migration Best Practices and Operational Guide*.

Minimally disruptive migration enables migrations on the same supported platforms as NDM it but requires a short application outage. The outage is because that the non-disruptive nature of migration is heavily dependent on the behavior of multipathing software to detect/enable/disable paths and is not in the control of Dell EMC (except for PowerPath).

NDM enables you to migrate storage group (application) data (the storage groups must have masking views) in a non-disruptive manner with no downtime for the following scenarios:

- From source arrays running Enginuity OS 5876 Q4 2016 and higher to target arrays running HYPERMAX OS 5977 Q4 2016 and higher.
- From source arrays running Enginuity OS 5876 Q4 2016 and higher to target arrays running PowerMax OS 5978.
- From source arrays running HYPERMAX OS 5977 Q3 2016 or higher to target arrays running PowerMax OS 5978.
- From source arrays running HYPERMAX OS 5977 Q3 2016 or higher to target arrays running HYPERMAX OS 5977 Q3 2016 and higher.
- From source arrays running PowerMax OS 5978 to target arrays running PowerMax OS 5978.

Source side service levels are automatically mapped to target side service levels.

NDM applies to open systems/FBA devices only.

NDM supports the ability to compress data on all-flash storage systems while migrating.

An NDM session can be created on a storage group containing session target volumes (R2s) where the SRDF mode is synchronous. The target volumes of an NDM session may also have a SRDF/Synchronous session that is added after the NDM session is in the cutover sync state.

#### Suggested best practices

- Try to migrate during slow processing times; QoS can be used to throttle copy rate.
- Use more SRDF links, if possible, to minimize impact:
  - Two is minimum number of SRDF links allowed; NDM can use up to eight SRDF links.
  - More links = more IOPS, lower response time.
- Use dedicated links as they yield more predictable performance than shared links.

You can migrate masked storage groups where the devices can also be in other storage groups. Examples of overlapping storage devices include:

- Storage groups with the exact same devices, for example, SG-A has devices X, Y, Z; SG-B has devices X, Y, Z.
- Overlapping devices, for example, SG-A has devices X, Y, Z ; SG-B has devices X, Y.
- Storage groups where there is overlap with one other migrated SG, for example, SG-A has devices X, Y, Z ; SG-B has devices W, X, Y ; SG-C has devices U,V,W.

The following migration tasks can be performed from Unisphere:

- [Setting up a migration environment](#) - Configures source and target array infrastructure for the migration process.
- [Viewing migration environments](#)
- [Creating a NDM session](#) - Duplicates the application storage environment from source array to target array.
- [Viewing NDM sessions](#)
- [Viewing NDM session details](#)
- [Cutting over a NDM session](#) - Switches the application data access from the source array to the target array and duplicates the application data on the source array to the target array.
- Optional: [Stop synchronizing data after NDM cutover](#) and [Start synchronizing data after NDM cutover](#) - stop or start the synchronization of writes to the target array back to source array. When stopped, the application runs on the target array only. Used for testing.

- Optional: [Cancelling a NDM session](#) - cancels a migration that has not yet been committed
- [Committing a NDM session](#) - Removes application resources from the source array and releases the resources that are used for migration. The application permanently runs on the target array.
- Optional: [Recovering a NDM session](#) - recovers a migration process following an error.
- [Removing a migration environment](#) - Removes the migration infrastructure.

## Understanding Virtual Witness

The Witness feature supports a third party that the two storage systems consult if they lose connectivity with each other, that is, their SRDF links go out of service. When this happens, the Witness helps to determine, for each SRDF/Metro Session, which of the storage systems should remain active (volumes continue to be read and write to hosts) and which goes inactive (volumes not accessible).

Prior to the HYPERMAX OS 5977 Q3 2016 or higher release, a Witness could only be a third storage system that the two storage systems involved in a SRDF/Metro Session could both connect to over their SRDF links.

The HYPERMAX OS 5977 Q3 2016 or higher release adds the ability for these storage systems to instead use a Virtual Witness (vWitness) running within a management virtual application (vApp) deployed by the customer.

The following Virtual Witness tasks can be performed from Unisphere.

[Viewing Virtual Witness instances](#)

[Adding a Virtual Witness](#)

[Viewing Virtual Witness instance details](#)

[Enabling a Virtual Witness](#)

[Disabling a Virtual Witness](#)

[Removing Virtual Witness](#)

## Understanding SRDF Delta Set Extension (DSE) pools

SRDF Delta Set Extension (DSE) pools provide a mechanism for augmenting the cache-based delta set buffering mechanism of SRDF/Asynchronous (SRDF/A) with a disk-based buffering ability.

This feature is useful when links are lost and the R1 system approaches the cache limitation. Data is moved out of cache into preconfigured storage pools set up to handle the excess SRDF/A data. When links recover, the data is moved back to cache and pushed over to the R2 system. DSE enables asynchronous replication operations to remain active when system cache resources are in danger of reaching system Write Pending (WP) or SRDF/A maximum cache limit.

## Understanding TimeFinder/Snap operations

TimeFinder/Snap operations enable you to create and manage copy sessions between a source volume and multiple virtual target volumes.

When you activate a virtual copy session, a point-in-time copy of the source volume is immediately available to its host through the corresponding virtual volume. Virtual volumes consume minimal physical disk storage because they contain only the address pointers to the data that is stored on the source volume or in a pool of SAVE volumes. SAVE volumes are storage volumes that are not

host-accessible and can only be accessed through the virtual volumes that point to them. SAVE volumes provide pooled physical storage for virtual volumes.

Snapping data to a virtual volume uses a copy-on-first-write technique. Upon a first write to the source volume during the copy session, Enginuity copies the preupdated image of the changed track to a SAVE volume and updates the track pointer on the virtual volume to point to the data on the SAVE volume.

The attached host views the point-in-time copy through virtual volume pointers to both the source volume and SAVE volume, for as long as the session remains active. If you terminate the copy session, the copy is lost, and the space that is associated with the session is freed and returned to the SAVE volume pool for future use.

**Note:** TimeFinder operations are not supported directly on storage systems running HYPERMAX OS 5977 or higher. Instead, they are mapped to their TimeFinder/SnapVX equivalents.

The following are the basic actions that are performed in a TimeFinder/Snap operation:

- **Create**—Creates the relationship between the source volume and the virtual target volume.
- **Activate**—Makes the virtual target volume available for read/write access and starts the copy-on-first-write mechanism.
- **Recreate**—Creates a point-in-time copy.
- **Restore**—Copies tracks from the virtual volume to the source volume or another volume.
- **Terminate**—Causes the target host to lose access to data pointed to by the virtual volume.

For more information about TimeFinder concepts, see the *Solutions Enabler TimeFinder Family CLI Product Guide* and the *TimeFinder Family Product Guide*.

## Understanding Open Replicator

Open Replicator is a non-disruptive migration and data mobility application.

When the Open Replicator control volumes are on a storage system running HYPERMAX OS 5977 or higher, the following session options cannot be used:

- Push
- Differential
- Precopy

There are many rules and limitations for running Open Replicator sessions. Refer to the *Solutions Enabler Migration CLI Product Guide* before creating a session. For a quick reference, see [Open Replicator session options](#).

## Open Replicator session options

Depending on the operation you are performing, some of the following options may not apply.

Session Option	Used with UI operation	Description
Consistent	Activate	Causes the volume pairs to be consistently activated.
	Donor Update Off	Consistently stops the donor update portion of a session and maintains the consistency of data on the remote volumes.

Session Option	Used with UI operation	Description
Copy	Create	Volume copy takes place in the background. This is the default for both pull and push sessions.
Cold	Create	Control volume is write disabled to the host while the copy operation is in progress. A cold copy session can be created as long as one or more directors discovers the remote device.
Differential	Create	Creates a one-time full volume copy. Only sessions created with the differential option can be recreated.  For push operations, this option is selected by default.  For pull operations, this option is cleared by default (no differential session).
Donor Update	Create	Causes data written to the control volume during a hot pull to also be written to the remote volume.
	Incremental Restore	Maintains a remote copy of any newly written data while the Open Replicator session is restoring.
Force	Terminate Restore	Select the <b>Force</b> option if the copy session is in progress. This will allow the session to continue to copy in its current mode without donor update.
	Donor Update Off	Select the <b>Force</b> option if the copy session is in progress. This will allow the session to continue to copy in its current mode without donor update.
Force Copy	Activate	Overrides any volume restrictions and allows a data copy.  For a push operation, remote capacity must be equal to or larger than the control volume extents and vice versa for a pull operation. The exception to this is when you have pushed data to a remote volume that is larger than the control volume, and you want to pull the data back, you can use the <b>Force_Copy</b> option.

Session Option	Used with UI operation	Description
Front-End Zero Detection	Create	Enables front end zero detection for thin control volumes in the session. Front end zero detection looks for incoming zero patterns from the remote volume, and instead of writing the incoming data of all zeros to the thin control volume, the group on the thin volume is de-allocated.
Hot	Create	Hot copying allows the control device to be read/write online to the host while the copy operation is in progress. All directors that have the local devices mapped are required to participate in the session. A hot copy session cannot be created unless all directors can discover the remote device.
Nocopy	Activate	Temporarily stops the background copying for a session by changing the state to CopyOnAccess or CopyOnWrite from CopyInProg.
Pull	Create	A pull operation copies data to the control device from the remote device.
Push	Create	A push operation copies data from the control volume to the remote volume.
Precopy	Create Recreate	For hot push sessions only, begins immediately copying data in the background before the session is activated.
SymForce	Terminate	Forces an operation on the volume pair including pairs that would be rejected. Use caution when checking this option because improper use may result in data loss.

## Understanding device groups

A device group is a user-defined group that consists of devices that belong to a locally attached array. Control operations can be performed on the group as a whole, or on the individual device pairs in the group. By default, a device can belong to more than one device group.

The user can create a legacy TF emulation from source devices with a SNAP VX snapshot. The prerequisites are:

- A SNAP VX storage group with a snapshot must exist.
- A device group must already have been created from this storage group.
- The device group must also have enough candidate target devices to create the required TF emulation session.

## Understanding SRDF groups

SRDF groups provide a collective data transfer path linking volumes of two separate storage systems. These communication and transfer paths are used to synchronize data between the R1 and R2 volume pairs that are associated with the SRDF group. At least one physical connection must exist between the two storage systems within the fabric topology.

See *Dell EMC SRDF Introduction* for an overview of SRDF.

The maximum number of supported SRDF groups differs by Enginuity version:

Enginuity	Maximum number of SRDF Groups supported			Group numbers
	per storage system	per director	per port	
5977 or higher	250	250	250	1–250
5876	250	64	64	1–250

- When specifying a local or remote director for a storage system running HYPERMAX OS 5977, you can select one or more SRDF ports.
- If the SRDF interaction includes a storage system running HYPERMAX OS 5977, then the other storage system must be running Enginuity OS 5876. Also, in this interaction the maximum storage system volume number that is allowed on the system running HYPERMAX OS 5977 is FFFF (65635).

## SRDF session modes

Mode	Description
Adaptive Copy	Allow the source (R1) volume and target (R2) volume to be out of synchronization by a number of I/Os that are defined by a skew value.
Adaptive Copy Disk Mode	Data is read from the disk and the unit of transfer across the SRDF link is the entire track. While less global memory is consumed it is typically slower to read data from disk than from global memory. Additionally, more bandwidth is used because the unit of transfer is the entire track. Additionally, because it is slower to read data from disk than global memory, device resynchronization time increases.
Adaptive Copy WP Mode	The unit of transfer across the SRDF link is the updated blocks rather than an entire track, resulting in more efficient use of SRDF

Mode	Description
	<p>link bandwidth. Data is read from global memory than from disk, thus improving overall system performance. However, the global memory is temporarily consumed by the data until it is transferred across the link.</p> <p>This mode requires that the device group containing the RDF pairs with R1 mirrors be on a storage system running Enginuity 5876.</p>
Synchronous	Provides the host access to the source (R1) volume on a write operation only after the storage system containing the target (R2) volume acknowledges that it has received and checked the data.
Asynchronous	<p>The storage system acknowledges all writes to the source (R1) volumes as if they were local devices. Host writes accumulate on the source (R1) side until the cycle time is reached and are then transferred to the target (R2) volume in one delta set. Write operations to the target device can be confirmed when the current SRDF/A cycle commits the data to disk by successfully de-staging it to the R2 storage volumes.</p> <p>For storage systems running Enginuity 5876, you can put an RDF relationship into Asynchronous mode when the R2 device is a snap source volume.</p>
AC Skew	Adaptive Copy Skew - sets the number of tracks per volume the source volume can be ahead of the target volume. Values are 0 - 65535.

## SRDF session options

Session option	Description	Available with action
Bypass	Bypasses the exclusive locks for the local and/or remote storage system during SRDF operations. Use this option only if you are sure that no other SRDF operation is in progress on the local and/or remote storage systems.	Establish Failback Failover Restore Incremental Restore Split Suspend Swap Write Disable R1



Session option	Description	Available with action
		Ready R1 Ready R2 RWDisableR2 Enable Disable
Consistent	Allows only consistent transition from async to sync mode.	Activate
Consistency Exempt	Allows you to add or remove volumes from an RDF group that is in Async mode without requiring other volumes in the group to be suspended.	Half Move Move Suspend
Establish	Fails over the volume pairs, performs a dynamic swap, and incrementally establishes the pairs. This option is not supported when volumes operating in Asynchronous mode are read/write on the RDF link. To perform a fail over operation on such volumes, specify the Restore option detailed higher in this table.	Failover
Force	Overrides any restrictions and forces the operation, even though one or more paired volumes may not be in the expected state. Use caution when checking this option because improper use may result in data loss.	Establish Incremental Establish Restore Incremental Restore Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable Disable Swap
Immediate	Causes the suspend, split, and failover actions on asynchronous volumes to happen immediately.	Suspend Split Failover
NoWD	No write disable - bypasses the check to ensure that the target of operation is write disabled to the host. This applies to the source (R1)	

Session option	Description	Available with action
	volumes when used with the Invalidate R1 option and to the target (R2) volumes when used with the Invalidate R2 option.	
SymForce	Forces an operation on the volume pair including pairs that would be rejected. Use caution when checking this option because improper use may result in data loss.	Restore Incremental Restore Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable Disable Swap
RecoverPoint Tag	Specifies that the operation will be performed on RecoverPoint volumes.	Restore Failback Refresh R1 Update
Refresh R1	Marks any changed tracks on the source (R1) volume to be refreshed from the target (R2) side.	Swap
Refresh R2	Marks any changed tracks on the target (R2) volume to be refreshed from the source (R1) side.	Swap
Remote	When performing a restore or failback action with the concurrent link up, data copied from the R2 to the R1 will also be copied to the concurrent R2. These actions require this option.	Restore Incremental Restore Failback
Restore	When the fail over swap completes, invalid tracks on the new R2 side (formerly the R1 side) will be restored to the new R1 side (formerly the R2 side).  When used together with the Immediate option, the fail over operation will immediately deactivate the SRDF/A session without waiting two cycle switches for session to terminate.	Failover

Session option	Description	Available with action
Star	<p>Selecting this option indicates that the volume pair is part of an SRDF/Star configuration. SRDF/Star environments are three-site disaster recovery solutions that use one of the following:</p> <ul style="list-style-type: none"> <li>• Concurrent SRDF sites with SRDF/Star</li> <li>• Cascaded SRDF sites with SRDF/Star</li> </ul> <p>This technology replicates data from a primary production (workload) site to both a nearby remote site and a distant remote site. Data is transferred in SRDF/Synchronous (SRDF/S) mode to the nearby remote site (referred to as the synchronous target site) and in SRDF/Asynchronous (SRDF/A) mode to the distant remote site (referred to as the asynchronous target site).</p> <p>SRDFR/Star is supported on Enginuity 5876. The <i>Solutions Enabler SRDF Family CLI Product Guide</i> contains more information on SRDF/Star.</p>	Establish Failback Failover Restore Incremental Restore Split Suspend Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable Disable

## SRDF/A control actions

Action	Activate Type	Write Pacing Type	Description
Activate	DSE	N/A	Activates the SRDF/A Delta Set Extension feature, which extends the available cache space by using device SAVE pools.

Action	Activate Type	Write Pacing Type	Description
	<b>Write Pacing</b> This feature extends the availability of SRDF/A by preventing conditions that result in cache overflow on both the R1 and R2 sides.	<b>Group write pacing</b> Group level write pacing is supported on storage systems running Enginuity 5876 and higher.	Activates SRDF/A write pacing at the group level.
		<b>Group &amp; Volume Write Pacing</b>	Activates SRDF/A write pacing at the group level and the volume level
		<b>Volume Write Pacing</b> Volume write pacing is supported on storage systems running Enginuity 5876 and higher.	Activates SRDF/A write pacing at the volume level.
	<b>Write Pacing Exempt</b>	N/A	Activates write pacing exempt. Write pacing exempt allows you to remove a volume from write pacing.

## SRDF group modes

The following values can be set for SRDF groups:

**Synchronous**—Provides the host access to the source (R1) volume on a write operation only after the storage system containing the target (R2) volume acknowledges that it has received and checked the data.

**Asynchronous**—The storage system acknowledges all writes to the source (R1) volumes as if they were local volumes. Host writes accumulate on the source (R1) side until the cycle time is reached and are then transferred to the target (R2) volume in one delta set. Write operations to the target volume can be confirmed when the current SRDF/A cycle commits the data to disk by successfully de-staging it to the R2 storage volumes.

For storage systems running Enginuity 5876, you can put an RDF relationship into Asynchronous mode when the R2 volume is a snap source volume.

**Semi Synchronous**—The storage system containing the source (R1) volume informs the host of successful completion of the write operation when it receives the data. The RDF (RA) director transfers each write to the target (R2) volume as the RDF links become available. The storage system containing the target (R2) volume checks and acknowledges receipt of each write.

**AC WP Mode On**—(adaptive copy write pending) the storage system acknowledges all writes to the source (R1) volume as if it was a local volume. The new data accumulates in cache until it is successfully written to the source (R1) volume and the remote director has transferred the write to the target (R2) volume.

**AC Disk Mode On**—For situations requiring the transfer of large amounts of data without loss of performance; use this mode to temporarily to transfer the bulk of your data to target (R2) volumes; then switch to synchronous or semi synchronous mode.

**Domino Mode On**—Ensures that the data on the source (R1) and target (R2) volumes are always in sync. The storage system forces the source (R1) volume to a Not Ready state to the host whenever it detects one side in a remotely mirrored pair is unavailable.

**Domino Mode Off**—The remotely mirrored volume continues processing I/Os with its host, even when an SRDF volume or link failure occurs.

**AC Mode Off**—Turns off the AC disk mode.

**AC Change Skew**—Modifies the adaptive copy skew threshold. When the skew threshold is exceeded, the remotely mirrored pair operates in the predetermined SRDF state (synchronous or semi-synchronous). As soon as the number of invalid tracks drop below this value, the remotely mirrored pair reverts back to the adaptive copy mode.

**(R2 NR If Invalid) On**—Sets the R2 device to Not Ready if there are invalid tracks.

**(R2 NR If Invalid) Of**—Turns off the (R2 NR\_If\_Invalid) On mode.

## RDFA flags

Flag	Status
(C)onsistency	X = Enabled, . = Disabled, - = N/A
(S)tatus	A = Active, I = Inactive, - = N/A
(R)DFA Mode	S = Single-session, M = MSC, - = N/A
(M)sc Cleanup	C = MSC Cleanup required, - = N/A
(T)ransmit Idle	X = Enabled, . = Disabled, - = N/A
(D)SE Status	A = Active, I = Inactive, - = N/A
DSE (A)utostart	X = Enabled, . = Disabled, - = N/A

## Understanding TimeFinder/Clone operations

Clone copy sessions enable you to create clone copies of a source volume on multiple target volumes.

The source and target volumes can be either standard volumes or BCVs, if they are the same size and emulation type (FBA/CKD). Once you have activated the session, the target host can instantly access the copy, even before the data is fully copied to the target volume.

**Note:** TimeFinder operations are not supported directly on storage systems running HYPERMAX OS 5977 or higher. Instead, they are mapped to their TimeFinder/SnapVX equivalents.

An overview of a typical clone session is:

1. Create a device group, or add volumes to an existing device group.
2. Create the session; restore the session.

3. Activate the session.
4. View the progress of the session.
5. Terminate the session.

For more information about TimeFinder/Clone concepts, see the *Solutions Enabler TimeFinder Family CLI Product Guide* and the *TimeFinder Family Product Guide*.

## Understanding TimeFinder/Mirror sessions

TimeFinder/Mirror is a business continuity solution that enables the use of special business continuance volume (BCV) devices. Copies of data from a standard device (which are online for regular I/O operations from the host) are sent and stored on BCV devices to mirror the primary data. Uses for the BCV copies can include backup, restore, decision support, and applications testing. Each BCV device has its own host address, and is configured as a stand-alone device.

- TimeFinder/Mirror requires Enginuity OS 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.
- TimeFinder operations are not supported on Open Replicator control volumes on storage systems running HYPERMAX OS 5977 or higher.

The TimeFinder/Mirror dashboard provides a single place to monitor and manage TimeFinder/Mirror sessions on a storage system.

## Understanding TimeFinder SnapVX

TimeFinder SnapVX is a local replication solution that is designed to nondisruptively create point-in-time copies (snapshots) of critical data.

TimeFinder SnapVX creates snapshots by storing changed tracks (deltas) directly in the Storage Resource Pool of the source volume. With TimeFinder SnapVX, you do not need to specify a target volume and source/target pairs when you create a snapshot. If the application ever needs to use the point-in-time data, you can create links from the snapshot to one or more target volumes. If there are multiple snapshots and the application needs to find a particular point-in-time copy for host access, you can link and relink until the correct snapshot is located.

## Understanding RecoverPoint

RecoverPoint provides block-level continuous data protection and continuous remote replication for on-demand protection and recovery at any point-in-time, and enables you to implement a single, unified solution to protect and/or replicate data across heterogeneous servers and storage.

RecoverPoint operations on Unisphere require Enginuity 5876 on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

## Understanding Performance Management

The Unisphere Performance Management application enables the user to gather, view, and analyze performance data to troubleshoot and optimize the storage systems.

Performance Management covers the following areas:

- Dashboards - Display predefined, user-defined custom dashboards, and templates.

- Charts - Create custom charts across multiple categories/metrics/time intervals.
- Analyze - Provide in-depth drill-down on storage system data for various collection ranges.
- Heatmap - Display hardware instances represented as colored squares, with the color indicating utilization levels.
- Reports - Create, manage, and run performance reports.
- Real-Time Traces - Create, manage, and run performance real-time traces .
- Databases - Manage Performance database tasks, for example, back up, restore, delete, as well as individual performance database information.
- Plan - Provide performance projection capacity dashboards displaying predicted future data that is based on linear projection.

## Understanding Database Storage Analyzer (DSA)

Database Storage Analyzer (DSA) is a feature that provides a database to storage performance troubleshooting solution for Oracle and MS SQL Server databases running on storage systems. No extra licensing cost is required in order to use DSA.

DSA is accessible by any Unisphere user. A DSA only user can also be created that has read only Unisphere access and can only view Databases section and Performance.

The main database list view presents I/O metrics such as response time, Input/Output Operations per second (IOPS) and throughput from both the database and the storage system which helps to immediately identify any gap between the database I/O performance and the storage I/O performance.

DSA offers the following benefits:

- Provides a unified view across database and storage.
- Quickly identifies when a database is suffering from high I/O response times.
- Reduces troubleshooting time for database and/or storage performance issues—DBAs and SAs can look at a unified database and storage I/O metrics view and quickly identify performance gaps or issues on both layers.
- Identifies database bottlenecks that are not related to the storage.
- Maps DB objects to storage devices
- Allows better coordination between the SA and DBA.
- Reduces repetitive manual drill downs for troubleshooting.

DSA supports the mapping of database files located on VMware virtual disks to their storage system volumes. With full database mapping, DSA can actively monitor 15-30 databases per Unisphere installation, depending on database size. Registering a database or instance with no extents mapping option allows the user to monitor hundreds of databases.

RAC and ASM are supported for Oracle, for CDB DSA guest user name should be started with c###. An Oracle diagnostic pack license is required for monitoring Oracle databases.

In addition, DSA supports FAST hinting capabilities for Oracle and MS SQL databases on storage systems running HYPERMAX OS 5977 or higher that allows users to accelerate mission-critical database processes in order to achieve improved response time. The user provides the timeframe, the database objects that should be hinted and the business priority. DSA then sends hints to the array in advance so that the FAST internal engine promotes those Logical Block Addresses (LBAs) to the right tier at the right time.

 **Note:** FAST hinting is only supported on hybrid arrays running HYPERMAX OS 5977 or higher.

## Understanding Unisphere support for VMware

Unisphere supports the discovery of vCenters or ESXi servers (using a read-only user) and integrates the information into Unisphere. VMware information is connected to its storage extents and this enables seamless investigation of any storage-related issues.

Unisphere support for VMware provides the storage admin access to all the relevant storage-related objects to an ESXi server and also helps troubleshooting storage performance-related issues to the ESXi server.

You can, as a read-only user, discover at the vCenter level and discover an individual ESXi server. If a vCenter is discovered, then all ESXi servers under that vCenter are discovered. All ESXi servers that do not have local storage on the Unisphere performing the discovery, are filtered out.

Once the user adds VMware information, all other users of Unisphere can access this information.

The minimum version number that vCenter supports is version 5.5.

## Understanding eNAS

Embedded NAS (eNAS) integrates the file-based storage capabilities of VNX arrays into storage systems running HYPERMAX OS 5977 or higher.

With this integrated storage solution, the Unisphere StorageAdmin provisions storage to eNAS data movers, which trigger the creation of storage pools in VNX. The users of Unisphere for VNX then use the storage pools for file-level provisioning, for example, creating file systems, file shares.

Unisphere provides the following features to support eNAS:

### **File System dashboard**

Provides a central location from which to monitor and manage integrated VNX file services.

### **Provision Storage for File wizard**

Allows you to provision storage to eNAS data movers.

### **Launch Unisphere for VNX**

Allows you to link and launch Unisphere for VNX.

## Understanding iSCSI

Unisphere provides monitoring and management for Internet Small Computer Systems Interface (iSCSI) directors, iSCSI ports, iSCSI targets, IP interfaces, and IP routes on storage systems running HYPERMAX OS 5977 or higher.

iSCSI is a protocol that uses the TCP to transport SCSI commands, enabling the use of the existing TCP/IP networking infrastructure as a SAN. As with SCSI over Fibre Channel (FC), iSCSI presents SCSI targets and devices to iSCSI initiators (requesters). Unlike NAS, which presents devices at the file level, iSCSI makes block devices available from the network. Block devices are presented across an IP network to your local system, and can be consumed in the same way as any other block storage device.

The iSCSI changes address the market needs originating from cloud/service provider space, where a slice of infrastructure, for example, computes, network and storage, is assigned to different users (tenants). Control and isolation of resources in this environment is achieved by the iSCSI changes. Also, more traditional IT enterprise environments also benefit from this new functionality. The changes also provide greater scalability and security.



## Understanding dynamic cache partitioning

Dynamic Cache Partitioning (DCP) divides the cache memory into multiple partitions with unique names and their device path assignments. Partition areas can be made static or dynamic in size.

The dynamic partitioning provides flexibility to the amount of floating memory that can be allocated with a high and low watermark. The flexibility enables memory resources to be temporarily donated to other partitions when needed. The `symqos` command enables you to create partitions for different device groupings besides the default partition that all devices belong to initially. Each partition has a target cache percentage and a minimum and maximum percentage. Also, you can donate unused cache to other partitions after a specified donation time.

**Note:** Enginuity OS 5876 is required for managing dynamic cache partitions. DCPs can be viewed on storage systems running HYPERMAX OS 5977 Q316SR or higher but they cannot be managed.

Copyright © 2019 Dell Inc. or its subsidiaries. All rights reserved.

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.