

EMC® Unisphere for VMAX

VERSION 1.6

Product Guide

REV 02

Copyright © 2011 - 2013 EMC Corporation. All rights reserved. Published in the USA.

Published May, 2013

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

CONTENTS

- Chapter 1: Getting Started with Unisphere for VMAX 19**
 - Operating as the initial setup user 20
 - Navigating the interface 21
 - Using the system selector 23
 - Discovering Symmetrix systems 24
 - Exporting data 24
 - Refreshing console information 24
 - Refreshing Symmetrix information 25
 - Setting preferences 25
 - Provisioning Symmetrix VMAX storage overview 26
 - Filtering lists of objects 28
 - Exiting the console 30
 - Getting help 30

- Chapter 2: Administration 31**
 - Setting system preferences 32
 - Displaying login messages 35
 - Alert settings 36
 - Runtime checks 36
 - Managing alert policies 37
 - Managing alert thresholds 37
 - Configuring alert notifications 39
 - Acknowledging alerts 41
 - Deleting alerts 41
 - Viewing alerts 42
 - Viewing alert details 43
 - Viewing alert thresholds 45
 - Viewing alert policies 46
 - Security 47
 - Authentication 47
 - Authorization Rules 50
 - Local Users 54
 - Link and launch 63
 - Creating link-and-launch client registrations 63

Editing link-and-launch client registrations	64
Deleting link-and-launch client registrations	64
Viewing link and launch client registrations	65
Chapter 3: System Management	67
Monitoring Symmetrix systems	68
System Dashboard:	68
Setting system attributes	70
Setting CPU I/O resource distribution	72
Audit log	73
Viewing the Symmetrix audit log	73
Viewing Symmetrix audit log details	73
Filtering audit log records	74
Hardware management	76
Converting directors	76
Setting director port attributes	77
Enabling and disabling director ports	79
Performing system health checks	79
Replacing failed drives	80
Locating Symmetrix VMAX systems	83
Viewing hardware components	84
Job management	97
Managing job lists	97
Making configuration changes safely	98
Understanding task persistence	98
Previewing/Running jobs	99
Scheduling jobs	100
Running jobs	100
Rescheduling jobs	100
Renaming jobs	100
Reordering tasks within a job	101
Grouping jobs	101
Un-grouping jobs	101
Stopping jobs	101
Deleting jobs	102
Viewing the job list	102
Viewing job details	104

Electronic licenses	106
Understanding licenses	106
Installing licenses	107
Removing host-based licenses	108
Viewing Symmetrix entitlements	108
Viewing host-based licenses	110
Viewing license usage	110
Host-based licenses	111
Symmetrix-based licenses	111
Symmetrix access controls	118
Understanding access controls	118
Opening access controls	118
Creating access groups	119
Adding access ID to access groups	119
Removing access IDs from access groups	120
Deleting access groups	120
Viewing access groups	120
Creating access pools	121
Modifying access pools	121
Deleting access pools	121
Viewing access pools	122
Modifying access types	122
Creating access control entries	123
Deleting access control entries	123
Viewing access control entries	123
Viewing access control entry details	124
Viewing access groups	124
Viewing access group details	125
Viewing access IDs	125
Viewing access pools	126
Viewing access pool details	126
Viewing access pool volumes	127
Viewing access types	127
Access types	128
Dynamic Cache Partitioning	130
Enabling/Disabling dynamic cache partitioning	130
Creating dynamic cache partitions	130

Assigning volumes to dynamic cache partitions	131
Deleting dynamic cache partitions	131
Running in analyze mode	131
Viewing dynamic cache partitions	132
Viewing dynamic cache partition details	133
Chapter 4: Storage Management	135
Assigning Symmetrix priority	136
Setting replication QoS	137
Symmetrix VMAX provisioning	139
Provisioning Symmetrix VMAX storage overview	139
Using the Provisioning Storage wizard	140
Storage groups	148
Creating storage groups	148
Adding volumes to empty storage groups	151
Expanding storage groups	152
Modifying storage groups	155
Associating/Disassociating FAST policies and storage groups	157
Adding or removing storage groups for cascaded storage groups	159
Cloning storage groups	159
Deleting storage groups	162
Setting host I/O limits	162
Managing VP compression on storage groups	163
Viewing storage groups	164
Viewing storage group details	166
Fully Automated Storage Tiering	169
Understanding FAST	169
Monitoring FAST	172
Symmetrix tiers	174
FAST policies	180
FASTcontroller	189
Time windows	191
Volume configuration	194
Creating regular volumes	194
Creating DATA volumes	195
Creating SAVE volumes	196
Creating thin volumes	197

Creating gatekeeper volumes	199
Creating diskless volumes	199
Creating DRV volumes	200
Creating volumes using storage templates	202
Deleting volumes	203
Duplicating volumes	203
Assigning Symmetrix priority	205
Changing volume configuration	205
Enabling and disabling volumes	206
Mapping volumes	206
Unmapping volumes	208
Setting optimized read miss	208
Setting volume status	210
Setting volume attributes	211
Setting volume identifiers	212
Setting volume names	212
Managing Meta Volumes	213
Disk groups	221
Renaming disk groups	221
Removing disks from disk groups	221
Deleting disk groups	221
Viewing disk groups	222
Viewing disk group details	223
Viewing disks in disk group	224
Viewing disk details	224
Virtual Provisioning	226
DATA volumes	226
Thin pools	231
Thin volumes	242
Enhanced Virtual LUN migration	248
Understanding Virtual LUN Migration	248
Migrating regular storage group volumes	248
Migrating regular volumes	249
Migrating thin storage group volumes	250
Migrating thin volumes	250
Terminating a VLUN migration session	251
Viewing VLUN migration sessions	251

Storage templates	252
Understanding storage templates	252
Creating storage templates	252
Modifying storage templates	254
Exporting storage templates	254
Importing storage templates	254
Deleting storage templates	255
Viewing storage templates	255
Viewing storage template details	256
External storage	258
Understanding external storage	258
Virtualizing external LUNs	258
Rescanning external storage	260
Viewing External Storage	261
Reservations	263
Reserving volumes	263
Adding volumes to reservations	263
Removing volumes from reservations	263
Releasing reservations	264
Viewing reservations	264
Viewing reservation details	264
Optimizer	266
Understanding Optimizer	266
Managing Optimizer	266
Starting/Stopping Optimizer	267
Enabling/Disabling Optimizer	268
Locking/Unlocking Optimizer	268
Approving Optimizer Swaps	268
Viewing Optimizer swap/move lists	269
Viewing Optimizer swap/move history	270
Chapter 5: Host Management	273
Initiators	274
Masking volumes	274
Unmasking volumes	275
Setting initiator port flags	275
Setting initiator attributes	275

Renaming initiator aliases	276
Replacing initiators	276
Removing masking entries	277
Viewing initiators	277
Viewing initiator details	278
Host aliases	281
Masking volumes	281
Creating host aliases	282
Adding initiators to host aliases	282
Removing initiators from host aliases	282
Unmasking volumes from host aliases	282
Deleting host aliases	283
Renaming host aliases	283
Viewing host aliases	283
Viewing host alias details	284
Masking views	285
Creating masking views	285
Renaming masking views	285
Deleting masking views	286
Viewing masking views	286
Viewing masking view connections	286
Viewing masking view details	288
Initiator groups	290
Creating initiator groups (Hosts)	290
Adding/Removing initiators from initiator groups	291
Adding/Removing initiator groups from initiator groups	291
Modifying initiator groups	291
Deleting initiator groups	293
Renaming initiator groups	293
Setting initiator group flags	293
Viewing initiator groups	294
Viewing initiator group details	294
Viewing initiators in initiator group	296
Port attributes	296
Port groups	298
Creating port groups	298
Deleting port groups	298

Adding ports to port group	299
Removing ports from port group	299
Renaming port groups	299
Viewing port groups	300
Viewing port groups details	300
Viewing ports in port group	301
Viewing port details	302
Host Cache Adapters	303
Viewing host cache adapters	303
Virtual servers	304
Adding a new virtual server	304
Adding storage to a VM	304
Removing a virtual server	305
Removing storage from a VM	305
Changing the password on a virtual server	306
Viewing virtual servers	307
Viewing the details of a virtual server	307
CU images	309
Mapping CKD volumes	309
Unmapping CKD volumes	309
Assigning alias addresses	310
Removing alias addresses	311
Assigning alias ranges	311
Removing alias ranges	312
Assigning alias counts	312
Removing alias counts	312
Viewing CU images	313
Viewing CU image details	314
Chapter 6: Data Protection	315
Local Replication	316
Monitoring local replication sessions	316
TimeFinder/Clone	319
TimeFinder/Snap	332
TimeFinder/Mirror	343
Device Groups	351
Creating device groups	351

Adding volumes to device groups	352
Removing volumes from device groups	352
Enabling device groups	352
Disabling device groups	352
Renaming device groups	353
Deleting device groups	353
Viewing device groups	354
Viewing device group details	354
Remote Replication	357
Managing remote replication sessions	357
SRDF	358
Properties panel	364
Related Objects panel	367
Migration	384
Understanding Virtual LUN Migration	384
Migrating regular storage group volumes	384
Migrating regular volumes	385
Migrating thin storage group volumes	386
Migrating thin volumes	386
Viewing VLUN migration sessions	387
Viewing VLUN migration session details	387
Replication Groups and Pools	389
SRDF/A DSE Pools	389
TimeFinder Snap Pools	392
SRDF Groups	395
RecoverPoint	409
Understanding RecoverPoint	409
Tagging and untagging volumes for RecoverPoint	410
Viewing RecoverPoint sessions	411
Viewing RecoverPoint session details	411
Viewing RecoverPoint tagged volumes	412
Viewing RecoverPoint tagged volume details	412
Open Replicator	415
Creating Open Replicator copy sessions	415
Managing Open Replicator sessions	415
Activating Open Replicator sessions	417
Restoring Open Replicator sessions	417

Renaming Open Replicator sessions	417
Removing Open Replicator sessions	418
Setting Open Replicator session background copy mode	418
Setting Open Replicator session donor update off	418
Setting Open Replicator session front end zero detection off	418
Setting Open Replicator session pace	419
Setting Open Replicator ceiling	419
Terminating Open Replicator sessions	419
Viewing Open Replicator sessions	419
Viewing Open Replicator session details	420
Open Replicator session options	421
Open Replicator flags	423
Federated Live Migration	425
Understanding Federated Live Migration	425
Setting up/Running Federated Live Migration	426
Creating a FLM session	428
Chapter 7: Performance	431
Using the Performance Viewer	432
Monitor view	433
Managing dashboards	433
Creating a dashboard folder	433
Creating a dashboard with charts	434
Creating a heatmap dashboard	434
Creating a dashboard for FAST	435
Editing a template dashboard	436
Copying a dashboard	436
Editing a dashboard	436
Deleting a dashboard	437
Viewing dashboards	438
Saving a dashboard as a template	438
Saving dashboard changes	439
Managing EMC dashboards	439
Using the FAST dashboards	440
Using the FAST by Tier dashboard	440
Using the FAST by Storage Group dashboard	441
Using the FAST by Policy dashboard	441

Scheduling a report from the dashboard	442
Running a report from the dashboard	442
Navigating from a heatmap to Analyze	442
Analyze view	444
Navigating in Analyze view	444
Monitoring data	444
About charts	445
Creating charts	448
Changing the time range	449
Creating a dashboard from Explore	450
Creating a template dashboard	450
Using the icons in Analyze view	451
Analyze Real Time	452
Analyze Diagnostic	454
Analyze Historical	472
Metric Tables	483
Settings	570
Exporting Performance Viewer settings	570
Exporting and importing Performance settings	570
System Registrations	571
Metrics	574
Thresholds and Alerts	575
Real Time Traces	577
Reports	579
Databases	584

FIGURES

Figure	Page
Figure 1: Chart styles	445
Figure 2: Data format	446
Figure 3: Display threshold	446

TABLES

Table	Page
Table 1: Job list configuration tasks	97
Table 2: Task status before and after server shutdown	98
Table 3: Symmetrix Licenses	111
Table 4: Comparison of FAST and Fast Virtual Pools (VP)	169
Table 5: Comparison of FAST DP and Fast VP	189
Table 6: Port attributes	296
Table 7: TimeFinder/Clone session options	330
Table 8: TimeFinder/Clone session options	342
Table 9: TimeFinder/Mirror session options	347
Table 11: Back-end director metrics	487
Table 12: Back-end director Real Time metrics	491
Table 13: Cache partition metrics	493
Table 14: Device pool metrics	495
Table 15: External disk group metrics	498
Table 16: External disk metrics	499
Table 17: Front-end ports metrics	501
Table 18: Device group and RDF group metrics	502
Table 19: Disk Buckets metrics	508
Table 20: Disk/Spare disk metrics	510
Table 21: Disk group metrics	512
Table 22: Disk group storage tier metrics	514
Table 23: FAST policy metrics	516
Table 24: Front-end director metrics	517
Table 25: Front-end director Real Time metrics	521
Table 26: Front-end ports metrics	524
Table 27: FE Port by SG metrics	525
Table 28: RDF director metrics	526
Table 29: RDF director Real Time metrics	528
Table 30: RDF/A group metrics	529
Table 31: Snap and DSE pool metrics	533
Table 32: Storage group explore metrics	534
Table 33: Symmetrix system Historical metrics	542
Table 34: Symmetrix system metrics	544
Table 35: Symmetrix system Real Time metrics	548

Table 36: Symmetrix Tier metrics	550
Table 37: Tier by Storage Group metrics	551
Table 38: Thin pool metrics	553
Table 39: Volume metrics	557
Table 40: Virtual pool tier metrics	565
Table 41: FE IO Limits by SG metrics	568
Table 42: FE Limits by FE metrics	569

CHAPTER 1

Getting Started with Unisphere for VMAX

This chapter contains the following:

- Operating as the initial setup user 20
- Navigating the interface 21
- Using the system selector 23
- Discovering Symmetrix systems 24
- Exporting data 24
- Refreshing console information 24
- Refreshing Symmetrix information 25
- Setting preferences 25
- Provisioning Symmetrix VMAX storage overview 26
- Filtering lists of objects 28
- Exiting the console 30
- Getting help 30

Operating as the initial setup user

When Unisphere for VMAX is first installed, there is a single user called the Initial Setup User (ISU). This user can perform administrative tasks only on Symmetrix systems that do not have defined roles (authorization rules). Once an Administrator or Security is assigned to a Symmetrix system, the ISU will no longer be able to access or even see the system from the Unisphere console. Therefore, it is recommended that users not operate in this role for too long.

The main tasks of an ISU are:

- ◆ [Creating local users](#) (page 56)
- ◆ [Authorization Rules](#) on page 50

Navigating the interface

The Unisphere web interface has seven major areas.

Section	Function
Home	<ul style="list-style-type: none"> ◆ View dashboard charts showing each system's usage. ◆ View and manage Administration Settings (lower right icon): <ul style="list-style-type: none"> • View and manage alert settings. • View and manage authentication options (local directory, Windows OS/AD, LDAP-SSL). • View and manage system preferences. • View and manage user authorizations. • View and manage link and launch client registrations.
System	<ul style="list-style-type: none"> ◆ View system dashboards. ◆ Display a list of active jobs ◆ Display a list of alerts. ◆ View the selected Symmetrix system's attributes. ◆ View and manage the visible Symmetrix systems' licenses.
Storage	<ul style="list-style-type: none"> ◆ View and manage storage groups (create, edit, delete, expand). ◆ View and manage the FAST™ controller policies. ◆ View and manage storage tiers. ◆ View and manage thin pools. ◆ View and manage storage volumes. ◆ View and manage storage templates. ◆ View external storage. ◆ View and manage disk groups.
Hosts	<ul style="list-style-type: none"> ◆ View and manage initiators. ◆ View and manage masking views. ◆ View and manage initiator groups (rename alias, set attributes, set flags, replace initiator). ◆ View a list of host aliases on the storage system.

Section	Function
	<ul style="list-style-type: none"> ◆ View and manage port groups.
Data Protection	<ul style="list-style-type: none"> ◆ View and manage local replication (create sessions, activate, recreate). ◆ Monitor and manage replication pools. ◆ Create and view device groups. ◆ Monitor and manage Symmetrix system migration sessions. ◆ Monitor and manage RecoverPoint™ sessions
Performance	<ul style="list-style-type: none"> ◆ Monitor and manage Symmetrix system dashboards (charts/graphs, heat maps, predefined for FAST). ◆ Perform trend analysis for future capacity planning. ◆ Analyze Symmetrix system data for diagnostic troubleshooting. ◆ Create charts for historical, diagnostic, and realtime Symmetrix system data. ◆ Manage policies for data collection and polling. ◆ Customize the performance metrics to your requirements. ◆ Set thresholds and alert notifications for Symmetrix system components. ◆ Maintain and schedule captures of a specified time period of realtime data. ◆ Create, schedule, execute, and export data queries. ◆ Configure the performance database and perform backups and restores.
Support	<ul style="list-style-type: none"> ◆ View online help for Unisphere tasks.

Using the system selector



The system selector, located at the far left of the navigation bar, displays a list of all managed Symmetrix systems. It allows you to toggle the interface between two modes, All Symmetrix and individual Symmetrix systems. Selections you make in the list provide the context for subsequent operations.

For example, to create a storage group on SYM00001234:

1. Select **SYM00001234** from the selector.
2. Select **Storage > Storage Groups** to open the **Storage Group** list view.
3. Click **Create** to open the **Create Storage Group** wizard.

Similarly, selections you make in the list will also be reflected in the interface's numerous views and lists. For example, if you are viewing a list of thin pools for SYM00001234 and you select SYM00005678, the view changes to list the thin pools for SYM00005678.

In addition, the system selector includes the **Discover Symmetrix** option that enables you to retrieve information on your Symmetrix storage environment. For more information, refer to [Discovering Symmetrix systems](#) (page 24).

Discovering Symmetrix systems

Discovery refers to process by which Symmetrix system, volume-level configuration and status information is retrieved. Discovered configuration and status data for all Symmetrix systems, as well as their directors and volumes, is maintained in a configuration database file on each host. Once your environment is discovered, you can direct information requests to retrieve Symmetrix-level (high-level) data or volume-level (low-level) information from the storage environment.


To discover a Symmetrix system:




1. Click the system selector, located at the far left of the navigation bar, and select **Discover Symmetrix**.
2. Click **OK** in the confirmation message.

Exporting data


This procedure explains how to export the contents of a view (list, details, dashboard, performance analyze) to a file:

1. While in the view, click export  in the title bar to open the **Export Wizard**.
2. Expand the **SMASApplicationContainer#** and select the component to export.
3. Click **Next**.
4. Select whether to export the content as **Data** in a table or as an **Image**, and click **Next**.
5. Select a format to use when exporting the data. The formats available here depend on whether you are exporting data or an image.
6. Click **Finish**.
7. Select a download location and click **Save**.


Refreshing console information

To refresh data in the Unisphere console window, click refresh  in the title bar. Unisphere for VMAX refreshes all its data from the database.

Refreshing Symmetrix information

To refresh Symmetrix system data, click Refresh Symmetrix  in the title bar. Unisphere for VMAX refreshes Symmetrix system data from the database. This operation will not discover new Symmetrix systems, only refresh data from existing systems.

Setting preferences

1. Click settings  in the title bar to open the **Preference Settings** dialog box.
2. Modify any of the following settings:
 - ◆ **Display Language** — Sets the language in which text is displayed in the interface. Only installed language packs are available.
 - ◆ **Logging Level** — Sets the severity level of the events that will display in the debug log.
 - ◆ **Optimize for Remote Connection** — Optimizes Unisphere for remote connections (for example, animations are removed). You should select this option when you are remotely connecting in a low-bandwidth situation.
3. Click **OK**.

Provisioning Symmetrix VMAX storage overview

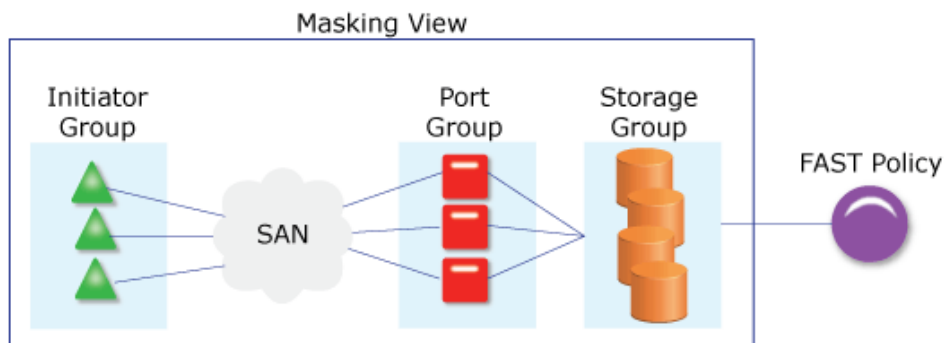
Provisioning storage refers to the process by which you make storage available to hosts.

EMC® Unisphere for VMAX provides the following methods for provisioning Symmetrix VMAX storage:

- ◆ **Recommended:** This method relies on wizards to step you through the provisioning process, and is best suited for novice and advanced users who do not require a high level of customization (i.e., the ability to create their own volumes, storage groups, etc.).
- ◆ **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 you will have a masking view, in which the volumes in the storage group are masked to the initiators in the initiator group and mapped to the ports in the port group.



Before you begin:

- ◆ The Symmetrix VMAX system must already be configured.
- ◆ For instructions on provisioning Symmetrix DMX storage, refer to [Masking volumes](#) (page 281).

To provision Symmetrix VMAX storage:

- | Recommended | Advanced |
|---|--|
| 1. Create a host (initiator group)
Use the Create Host Wizard to group host initiators (HBAs). | 1. Create a host (initiator group)
Use the Create Host Wizard to group host initiators (HBAs) |
| 2. Provision storage for the new host
Use the Provision Storage Wizard, which will | 2. Create the volumes
Create one or more volumes on the Symmetrix system. |

step you through the process of creating the storage group, port group, and masking view, and to optionally associate the storage group with a FAST policy.

3. [Create a storage group](#)
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 wizard's **Storage Group Type** to **Empty**, and then complete [Adding volumes to empty storage groups](#).
4. [Create a port group](#)
Group Fibre Channel and/or iSCSI front-end directors.
5. [Create a masking view](#)
Associate the initiator, storage, and port groups into a masking view.
6. [Associate the storage group with a FAST policy](#)
Optional: Associate the storage you created in step 3 with an existing FAST policy and assign a priority value for the association.


Filtering lists of objects

Many list views, wizards, and dialog boxes include a filter tool to help you quickly locate specific objects.


Unisphere for VMAX includes two types of filtering tools to help you quickly locate specific objects in a list:

- ◆ **Simple** — Allows you search for specific objects by name. For example, the name of a storage group.
- ◆ **Advanced** — Allows you to filter lists by attributes. For example, capacity, reservations, etc. This type of filter is only available on volume list views.

To use the simple filter:

1. Click the filter icon  to display a search field.
2. Type all or part of the object's name and press **Enter**.
3. To clear the filter (that is, view the complete list again), click the filter icon again and select **Clear**.

To use the advanced filter

1. Click the filter icon  to open the **Advanced Filter** dialog box.
2. Select or type a value for any number of the following criteria, and then click **OK**.
 - Volume Type**
 - ◆ **Capacity (MB)** — Filters the list for volumes with a specific capacity.
 - ◆ **Capacity Range** — Filters the list for volumes with capacities within the range.
 - ◆ **Volume Configuration** — Filters the list for volumes with a specific configuration.
 - ◆ **Advanced:**
 - **Emulation** — Filters the list for volumes with a specific emulation.
 - **Private Volumes** — Filters the list for volumes of a specific type.
 - **Meta** — Filters the list for metavolumes of a particular type.
 - **Volume Flag** — Specifies any volume flags.
 - Volume Identifier**
 - ◆ **Volume ID** — Filters the list for a volume with specific ID.
 - ◆ **Volume Range** — Filters the list for volumes with IDs within the range.
 - ◆ **Volume Identifier Name** — Filters the list for the specified volume name.
 - ◆ **Advanced:**
 - **Volume Nice Name** — Filters the list for a volume with a specific nice name.
 - Volume Availability**
 - ◆ **Status** — Filters the list for volumes with a specific status.

- ◆ **Reservation** — Filters the list for reserved volumes of a specific type.
- ◆ **Used:**
 - **Mapped** — Specifies whether to include/exclude mapped volumes.
 - **Masked** — Specifies whether to include/exclude masked volumes.
 - **Bound** — Specifies whether to include/exclude bound volumes.
 - **Enabled** — Specifies whether to include/exclude enabled volumes.
 - **Held** — Specifies whether to include/exclude held volumes.
 - **In Storage Group** — Specifies whether to include/exclude volumes that are in storage groups.
 - **In Device Group** — Specifies whether to include/exclude volumes that are in device groups.
 - **AS400 Gatekeeper** — Specifies whether to include/exclude AS400 gatekeeper volumes.
 - **Host Cache Card** — Specifies whether to include/exclude volumes under host cache card control.
 - **D1F1** — Specifies whether to include/exclude volumes that do not have the d1f1 attribute set.

Replication

- ◆ **SRDF group** — Filters the list for volumes within a specific SRDF group.
- ◆ **Dynamic RDF** — Filters the list for dynamic SRDF volumes of a specific type.
- ◆ **Used:**
 - **RecoverPoint** — Specifies whether to include/exclude volumes under RecoverPoint control.
 - **Write Pacing Capable** — Specifies whether to include/exclude volumes that are capable of using SRDF/A write pacing.
 - **Concurrent RDF** — Specifies whether to include/exclude volumes with the Concurrent RDF flag set.
 - **Cascaded RDF** — Specifies whether to include/exclude volumes with the Cascaded RDF flag set.
 - **Diskless RDF** — Specifies whether to include/exclude volumes with the Diskless RDF flag set.
 - **RDF not ready if Invalid** — Specifies whether to include/exclude volumes with the RDF_NR_IF_INV flag set.
 - **RDF Asynchronous** — Specifies whether to include/exclude volumes with the RDF Asynchronous flag set.
 - **Attached Target** — Specifies whether to include/exclude volumes with the Attached Target flag set.
 - **Attached BCV** — Specifies whether to include/exclude volumes with the Attached BCV flag set.

Related Objects

- ◆ **Disk Technology** — Filters the list for volumes on a specific disk technology.
- ◆ **Pool** — Filters the list for volumes in a specific pool.
- ◆ **Disk Group** — Filters the list for volumes in a specific disk group.
- ◆ **Storage Group** — Filters the list for volumes in a specific storage group.
- ◆ **Device Group** — Filters the list for volumes in a specific device group.
- ◆ **Tiers** — Filters the list for volumes in a specific Symmetrix tier.

Federated Tiered Storage

- ◆ **Used:**
 - **External** — Specifies whether to include/exclude external volumes.
 - **Encapsulated** — Specifies whether to include/exclude encapsulated volumes.
 - **Encapsulated WWN** — Specifies whether to include/exclude volumes with encapsulated WWNs.
 - **Geometry Limited** — Specifies whether to include/exclude volumes with limited geometry.

Virtual Provisioning


- ◆ **Allocated Capacity Percentage** — Filters the list for volumes with a specific allocated capacity.
- ◆ **Written Capacity Percentage** — Filters the list for volumes with a specific written capacity.
- ◆ **Bound to Thin Pool** — Filters the list for volumes bound to specific thin pools.

1. To clear the filter, open the **Advanced Filter** dialog box, click **Clear All**, and then **OK**.

Exiting the console

To exit the Unisphere for VMAX console, click exit  in the title bar.

Getting help

Clicking help  in the title bar opens the entire help system.

Clicking help in a dialog box, wizard page, or view opens a help topic specifically for that dialog, page, or view.

CHAPTER 2

Administration

This chapter contains the following:

Setting system preferences	32
Displaying login messages	35
Alert settings	36
Security	47
Link and launch	63

Setting system preferences

This procedure explains how to set preferences for a Symmetrix system.

Before you begin:

Only a user with Administrator permission can set system preferences.

To set system preferences:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Preferences** to open the **Preferences** page.

The following preferences can be set or default values changed:

- ◆ **Debug** — Specify debug level. Set the following parameters:
 - **Debug** — Set the level of debugging to write to the debug file.
 - **Debug2** — Set the secondary level of debugging to write to the debug file.
 - **Debug Filename** — Enter the debug file name.
- ◆ **Reservation** — Specify Symmetrix system volume reservation preferences. Set the following parameters:

Reservations must be enabled and enforced in the Options file which is located in the SYMAPI configuration directory. The settings are displayed in the **Reservation** panel.

- **Enable Reservation** — Specifies whether volume reservations are enabled for this host.
- **Reservation Type** — Specifies whether a volume reservation is enforced for all users.
- **Default Reservation after expiration days** — Enter the number of days that volume reservations will remain in affect. The value specified here will appear as the default expiration date in all dialog boxes and wizard panels that include the **Reserve** option. Setting this value to 0 (default) disables this option and the **Reserve** option will appear clear (unchecked) in all dialogs and wizard panels.

Select to enable or clear to disable for the following configuration parameters:

- **Enable acquire upon job start** — Allows for automatic volume reservation on volumes for a job.
- **Enable require comments upon acquire** — Specifies that the user must enter comments for volume reservations. When enabled, the user must enter a comment, or an error displays.
- **Enable auto release after job completion** — Allows the volume reservation to be released after the job has been committed.
- **Enable auto release after clear** — Allows the volume reservation to be released after the task has been cleared from the job list.
- ◆ **Replication** — Specify the following Replication parameters:

- **BCV Delay (sec)** — Select the amount of time to wait between TimeFinder® establish operations. The delay can be set from 0 to 30 seconds. The default value is 0.
- **BCV Establish Type** — Select the TimeFinder establish type.
- **BCV Cancel Policy** — Select the BCV pair cancel policy for establishing new pairs when the maximum number of pairs has been reached.
- **Max BCV Pairs** — Select the maximum number of BCV pairs (0-16).
- **Clone Copy Mode** — Select the default behavior for creating clone sessions.
- **Clone Pair Policy** — Select the clone terminate policy when establishing a new clone, and the maximum number of clones has been reached.
- **Automatic Clone Creation** — Select the mode in which to create the clone sessions. Possible values are:
 - **Copy No Diff** — Creates a nondifferential (full) copy session in the background.
 - **PreCopy No Diff** — Creates a nondifferential (full) copy session in the background before the activate starts.
 - **Copy Diff** — Creates a differential copy session in the background. In differential copy sessions, only those volume tracks that have changed since the full clone was performed are copied (that is, only new writes to the source volume will be copied).
 - **PreCopy Diff** — Creates a differential copy session in the background before the activate starts. In differential copy sessions, only those volume tracks that have changed since the full clone was performed are copied (that is, only new writes to the source volume will be copied).
 - **VSE No Diff** — Creates a VP Snap Copy session.
- **Clone Target** —
- **Enable clone copy on write** — Restricts copy operations when the target of a clone is being read. Select to enable, or clear to disable.
- **Enable clone larger target** — Allows cloning from a source volume to a larger target volume. Select to enable, or deselect to disable.
- **Enable command scope** — Limits eligible volumes for TimeFinder operations to source volumes that have sessions with target volumes contained within the same device group or composite group. Enabling this option limits the eligible volumes. This option is only available on Symmetrix systems running Enginuity 5874 or higher. Select to enable, or clear to disable.
- **Enable multi-virtual snap** — Allows for up to 128 snap sessions on the same source volume. For Symmetrix systems running Enginuity 5874.207.166 or higher, this also allows for the creation of multi-virtual snap sessions from thin volumes. Select to enable, or clear to disable.
- **Open Replicator Copy Mode** — Select the default behavior for creating Open Replicator copy sessions.
- **Snap Pair Policy** — Select the snap terminate policy for establishing a new snap when the maximum number of snaps has been reached.

- **Enable RDF group-level consistency** — Allows for checks for consistency of the RDF group level during a consistent operation. Select to enable, or clear to disable.
 - **Enable TF/Clone emulation** — Allows for TF/Clone emulation. Select to enable, or clear to disable.
 - ◆ **Login Message** — Type a message to display to EMC® Unisphere for VMAX users during login. For example, an you may want to notify logging in users about a software upgrade. Messages can be up to 240 characters/6 lines.
4. Click **Apply**.

Displaying login messages

The login message feature enables Administrators and StorageAdmins to display a message to Unisphere for VMAX users during login. For example, an administrator may want to notify logging in users about a software upgrade.

To create a login message:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Preferences** to open the **Preferences** page.
4. In the **Login Message** panel, type the message you want to display. Messages can be up to 240 characters.

Alert settings

You can configure Unisphere for VMAX to monitor Symmetrix systems for specific events or error conditions. When an event/error of interest occurs, Unisphere for VMAX will display an alert and, if configured to do so, notify you of the alert by way of email, SNMP, or Syslog.

In addition to alerting you of specific events/errors, Unisphere for VMAX also performs a number of runtime checks, for which it will also alert you. For more information, refer to [Runtime checks](#) below.

The procedures in this section explain how to configure and use the alert functionality.

Runtime checks

Unisphere for VMAX performs the following runtime checks. If a check returns a result outside the expected range, Unisphere for VMAX will raise the appropriate alert.

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

Runtime check	Threshold	Alert
Total memory on the Unisphere server	Unisphere = 4 GB* Unisphere and Performance = 8 GB	System memory (# GB) is below the minimum requirement of (# GB)
Free disk space on the Unisphere installed directory	20 GB	Free disk space <# GB> is below the minimum requirement of 20 GB
Number of managed Symmetrix systems	10	Number of managed arrays (#) is over the maximum supported number of 10
Number of managed volumes	80,000	Number of managed volumes (#) is over the maximum supported number of 80,000. Note that Solutions Enabler may indicate a slightly different number of volumes than indicated in this alert.
Number of gatekeepers	6	Number of gatekeepers (#) on Symmetrix (<i>SymmID</i>) is below the minimum requirement of 6.

* If running Unisphere for VMAX on the Symmetrix service processor, this threshold is 2 GB.

** If running Unisphere for VMAX on the Symmetrix service processor, this threshold is one gatekeeper.

Managing alert policies

This procedure explains how to enable/disable the policies for which Unisphere for VMAX monitors and reports alerts.

Before you begin:

- ◆ To perform this operation, you must be an Administrator or StorageAdmin.
- ◆ To receive alert notifications, you must first [configure the alert notifications feature](#).
- ◆ For alert (event) descriptions, refer to the *EMC Solutions Enabler Installation Guide*.

To enable/disable alert policies:

1. From the **System Selector**, select **All Symmetrix**.
2. In the **Common Tasks** panel, click **Administration > Alert Settings > Alert Policies** to open the **Alert Policies** list view.
3. Do the following, depending on whether you are enabling or disabling an alert:

Enabling:

- a. Select one or more alerts, and click **Enable**.
- b. To enable EMC® Unisphere for VMAX to notify you when a Symmetrix system reports an alert, click **Notify** to open the **Define Notification** dialog box.
- c. Select the **Symmetrix** system on which to define the notification.
- d. Select the **Policy Names** for which you want Unisphere to notify you.
- e. Select one or more **Notification Types**.
- f. Click **OK** to close the **Define Notification** dialog box.

Disabling:

Select one or more alerts, and click **Disable**.

Managing alert thresholds

Certain alerts are associated with a numerical value. This value is compared with a set of threshold values, which determine whether the alert is delivered and, if so, with what severity. This procedure explains how to manage the alert threshold feature.

Before you begin:

- ◆ Pool utilization thresholds are enabled by default on every Symmetrix VMAX Family system.
- ◆ To receive utilization threshold alerts, you must [enable alerts](#) on the Symmetrix system.
- ◆ To receive alert notifications, you must first [configure the alert notifications feature](#).

To manage alert thresholds:

1. From the **System Selector**, select **All Symmetrix**.

2. Click **Administration** (in the **Common Tasks** panel) › **Alert Settings** › **Alert Thresholds** to open the **Alert Thresholds** list view.
3. Do the following, depending on whether you are creating, editing, or deleting thresholds:

Creating:

- a. Click **Create** to open the **Create Threshold Policies** dialog box.
- b. Select the **Symmetrix** system on which to create the threshold.
- c. Select the **Category** of threshold to assign. Possible values are:
 - ◆ **DSE Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a DSE pool.
 - ◆ **DSE Pool Spillover** — Threshold event that reflects how long (in minutes) an SRDF spillover has been occurring.
 - ◆ **Snap.Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a snap pool.
 - ◆ **Thin Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a virtual pool.
 - ◆ **FAST VP Policy Utilization** — Threshold event that reflects the allocated capacity (as percentage) of all the pools in all the tiers in a FAST VP policy.
- d. Select the pools (**Instances to enable**) on which to create the threshold.
- e. Specify a threshold value (percentage of utilization) for each severity level: **Warning**, **Critical**, and **Fatal**.
- f. Click **OK** to close the **Create Threshold Policy** dialog box.

Editing:

- a. Select a threshold and click **Edit** to open the **Edit Threshold Policies** dialog box.
- b. Specify a new threshold value (percentage of utilization) for any number of the severity levels: **Warning**, **Critical**, and **Fatal**.
- c. Click **OK** to close the **Edit Threshold Policy** dialog box.

Deleting:

Select one or more thresholds and click **Delete**.

Configuring alert notifications



The alert notification feature allows you to configure EMC® Unisphere for VMAX to deliver alert notifications.

Before you begin:

- ◆ To perform this operation, you must be an Administrator or StorageAdmin.
- ◆ EMC® Unisphere for VMAX employs the following throttling algorithms to prevent alert flurries from straining the system.
 - Symmetrix Event Throttling — When a Symmetrix system raises an alert flurry, the alert infrastructure packages all the alerts into a single notification.
 - Generic Throttling — When the number of alerts generated by a non-Symmetrix event exceeds a set threshold, the alert infrastructure will ignore subsequent alerts from the source.

To configure alert notifications:

1. From the **System Selector**, select **All Symmetrix**.
2. Click **Administration** (in the **Common Tasks** panel) > **Alert Settings** > **Notifications** to open the **Notifications** page.
3. Specify the method for delivering the alert notifications (**Notification Types**):
 - a. Select or clear any number of the following:
 - ◆ **Enable Email** — Forwards alert notifications to an e-mail address.
 - ◆ **Enable SNMP** — Forwards alert notifications to a remote SNMP listener.
 - ◆ **Enable Syslog** — Forwards alert notifications to a remote syslog server.

- b. If enabling a new notification type, notice whether you need to configure a server. A green checkmark () indicates that server is configured. A red X () indicates that the server is not configured, in which case you will need to complete the corresponding procedures below.
 - c. Click **Apply**.
4. Specify the alert levels for which you want to receive notifications **Alert Levels for Notifications:**
 - To edit alert levels:**
 - a. Select the Symmetrix system and click **Edit** to open the **Edit Alert Level for Notification** dialog box.
 - b. Select/clear the alert levels and click **OK**. Note that Performance Alerts are only available if you have the Performance option installed.
 - c. Repeat step 4 for each Symmetrix system.
 - To remove all alert levels:**

Select the Symmetrix system and click **Delete**.
5. Complete any number of the following procedures, depending on the notification types you selected, and whether you need to configure the corresponding server.

Configuring an Email server:

1. Click **Config Email Server** to open the **Configure Email Server** dialog box.
2. Type the email address of the email server (**Sender Email Address**).
3. Type the IP address of the email server (**IP Address/Host**).
4. If the port of you email server differs from the default (25), type a new value (**Server Port**).
5. *Optional:* To send a test message to the email server, click **Test**.
6. Click **OK** to close the **Config Email Server** dialog box.
7. Specify receiver addresses:

To add an address:

- a. Select the Symmetrix system for which you want to receive notifications and click **Edit** to open the **Edit Email Address** dialog box.
- b. To receive Array/System alert notifications, select **SMC** and type the recipient's email address.
- c. To receive Performance alert notifications, select **SPA** and type the recipient's email address.
- d. Click **OK** to close the **Edit Email Address** dialog box.

To delete an address:

Select the Symmetrix system and click **Delete**.

Configuring an SNMP server:

To add an SNMP server:

1. Click **Add/Edit** to open the **Add/Edit SNMP Server** dialog box.
2. Type IP address of the SNMP server (**IP Address/Host**).
3. Type the port on the SNMP sever (**Server Port**).
4. Click **OK** to close the **Add SNMP Server** dialog box.

To verify that traps are configured correctly:

Select the server and click **Test**, and then verify that you received a trap.

To delete an SNMP server:

Select the server and click **Remove**.

Configuring a remote Syslog server:

Refer to *Setting up the event daemon for monitoring* in the *Solutions Enabler Installation Guide* for instructions.

Acknowledging alerts

1. Select the Symmetrix system.
2. Select **System** › **Alerts** to open the **Alerts** list view.
3. Select one or more alerts and click **Acknowledge**.

Deleting alerts

The following explains how to delete System and Array alerts:

1. Select the Symmetrix system.
2. Select **System** › **Alerts** to open the **Alert** list view.
3. Select one or more alerts and click **Delete**.

Viewing alerts

This procedure explains how to view alerts for a particular Symmetrix system or all the visible Symmetrix systems.

Before you begin:

- ◆ For alert (event) descriptions, refer to the *EMC Solutions Enabler Installation Guide*.
- ◆ In addition to alerting you of specific events/errors, Unisphere for VMAX also performs a number of runtime checks, for which it will also alert you. For more information, refer to [Runtime checks](#) on page 36.

To view alerts:

1. Do the following, depending on whether you want to view the alerts for a particular Symmetrix system, or all Symmetrix systems.

For a particular Symmetrix system:

- a. Select the Symmetrix system.
- b. Select **System** › **Alerts** to open the system's **Alerts** list view.

For all visible Symmetrix systems:

- a. Select **All Symmetrix** systems.
- b. At the bottom of the **Home** section, click **Alerts** to open the **All Alerts** view.

The following properties display:

- ◆ **State** — State of the alert. Possible values are New or Acknowledged.
- ◆ **Severity** — Alert's severity. Possible values are:
 - (1) Fatal
 - (2) Critical
 - (3) Warning
 - (4) Information
 - (5) Normal
- ◆ **Type** — Type of alert. Possible values are Array, Performance, Server, and System.
- ◆ **Symmetrix** — Symmetrix system reporting the alert. This field only appears when viewing alerts for all Symmetrix systems. This field will appear blank for server alerts. This is because server alerts are specific to the server or runtime environment and are not associated with a specific object or Symmetrix system.
- ◆ **Object** — Object to which the alert is related. This field will appear blank for server alerts. This is because server alerts are specific to the server or runtime environment and are not associated with a specific object or Symmetrix system.
- ◆ **Description** — Description of the alert.
- ◆ **Created** — Date/time the alert was created.
- ◆ **Acknowledged** — Date/time the alert was acknowledged.
- ◆ **Acknowledger** — ID of the user who acknowledged the alert.

The following controls are available:

- ◆ **View Details** — See [Viewing alert details](#) below.
- ◆ **Acknowledge** — See [Acknowledging alerts](#) on page 41.
- ◆ **Delete** — See [Deleting alerts](#) on page 41.

Viewing alert details

1. Select the Symmetrix system.
2. Select **System** › **Alerts** to open the **Alert** list view.
3. Select an alert and click **View Details** to open its **Details** view.

Use the Symmetrix alert **Details** view to display and manage a Symmetrix alert. This view contains two panels, [Properties](#) and [Related Alerts](#) or [Causes](#).

Properties panel

The following properties display:

- ◆ **Alert ID** — Unique number assigned by EMC® Unisphere for VMAX.
- ◆ **State** — State of the alert. Possible values are new or acknowledged.
- ◆ **Severity** — Alert's severity. Possible values are:
 - (1) Fatal
 - (2) Critical
 - (3) Warning
 - (4) Information
 - (5) Normal
- ◆ **Type** — Type of alert. Possible values are Array, Performance, and System.
- ◆ **Symmetrix** — ID of the Symmetrix system generating the alert.
- ◆ **Object** — Object to which the alert is related. For more information, click the object to open its details view.
- ◆ **Metric** — Key performance index. This field only displays for Performance alerts.
- ◆ **Value** — Metric's data. This field only displays for Performance alerts.
- ◆ **Count** — How many times the metric and its value are exceeded. This field only displays for Performance alerts.
- ◆ **Created** — Date/time the alert was created.
- ◆ **Description** — Description of the alert.

The following controls are available:

- ◆ **Acknowledge** — See [Acknowledging alerts](#) on page 41.
- ◆ **Delete** — See [Deleting alerts](#) on page 41.

Related Alerts panel

The **Related Alerts** panel links you to other alerts related to the same object, if any. This panel only displays for Array alerts.

Causes panel

The **Causes** panel links you to possible causes of the alert. This panel only displays for Performance alerts.

Viewing alert thresholds

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home** > **Administration** (in the **Common Tasks** panel) > **Alert Settings** > **Alert Thresholds** to open the **Alert Thresholds** list view.

Use the **Alerts Thresholds** list view to display and manage alert thresholds.

The following properties display:

- ◆ **Symmetrix ID** — Symmetrix system on which the threshold is defined.
- ◆ **Category** — Category on which the threshold is defined. Possible values are:
 - **DSE Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a DSE pool.
 - **DSE Pool Spillover** — Threshold event that reflects how long (in minutes) an SRDF spillover has been occurring.
 - **Snap.Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a snap pool.
 - **Thin Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a virtual pool.
 - **FAST VP Policy Utilization** — Threshold event that reflects the allocated capacity (as percentage) of all the pools in all the tiers in a FAST VP policy.
- ◆ **Instance** — Pool on which the threshold is defined. An asterisk (*) indicates that the default pool type was used.
- ◆ **State** — Whether alert thresholds are enabled or disabled for the pool.
- ◆ **Notification** — Whether the alert notification option is enabled or disabled for the alert.
- ◆ **Warning** — Percentage of utilization at which point a warning alert will be issued.
- ◆ **Critical** — Percentage of utilization at which point a critical alert will be issued.
- ◆ **Fatal** — Percentage of utilization at which point a fatal alert will be issued.
- ◆ **Custom** — Whether the policy has been customized.

The following controls are available:

- ◆ **Create** — See [Managing alert thresholds](#) on page 37.
- ◆ **Edit** — See [Managing alert thresholds](#) on page 37.
- ◆ **Notify** — See [Configuring alert notifications](#) on page 39.
- ◆ **Delete** — See [Managing alert thresholds](#) on page 37.

Viewing alert policies

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home** > **Administration** (in the **Common Tasks** panel) > **Alert Settings** > **Alert Policies** to open the **Alert Policies** list view.

Use the **Alerts Policies** list view to display and manage alert policies.

The following properties display:

- ◆ **Symmetrix ID** — Symmetrix system on which the policy is defined (for Array type policies) and an asterisk (*) for SMAS type policies.
- ◆ **Type** — Type of alert policy. Possible values are **Array** for Symmetrix-based alerts and **SMAS** for application-based alerts.
- ◆ **Policy Name** — Name of the policy. For alert (event) descriptions, refer to the *EMC Solutions Enabler Installation Guide*.
- ◆ **State** — Whether the policy is **Enabled** or **Disabled**.
- ◆ **Notification** — Icon indicating the method to use when delivering the alert notification (e-mail, SNMP, or Sys Log). A blank indicates that Unisphere is not configured to deliver an alert notification for the corresponding policy.

The following controls are available:

- ◆ **Enable** — See [Managing alert policies](#) on page 37.
- ◆ **Disable** — See [Managing alert policies](#) on page 37.
- ◆ **Notify** — See [Configuring alert notifications](#) on page 39.

Security

Authentication

Login authentication

When you log in, Unisphere for VMAX 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.)

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

Configuring authentication authorities

This procedure explains how to configure EMC® Unisphere for VMAX to authenticate users.

Before you begin:

If configuring authentication to use LDAP, obtain the LDAP-SSL server bind distinguished name (DN) and password from your LDAP Administrator.

To configure authentication:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Authentication** to open the **Authentication** list view.
5. Click **Edit** to open the **Configure Authentication** wizard.
6. Select the **Authentication Authority** to use during login. Possible values are:
 - ◆ **Local Directory** — When enabled, users can log in as a user from the CST local directory.
 - ◆ **LDAP-SSL** — When enabled, users can log in as a user from the configured LDAP directory.
 - ◆ **Windows OS/AD** — When enabled, users can log in as a user from the Windows local host and/or from the Active Directory domain. This option only applies to Windows installations.

7. If you selected the Windows OS/AD authority, as an option you can specify to limit authentication to members of a specific Windows OS/AD group. To do this, click **Show Advanced**, select the option, and type the **Group Names**, separated by commas.
8. Click **Next**.
9. If you are configuring LDAP-SSL, do the following:
 - a. Specify values for the following parameters and click **Next**.
 - ◆ **Server** — IP address or hostname of the LDAP server to use for authentication. Only alphanumeric characters are allowed. Values longer than 40 characters will wrap.
 - ◆ **Port** — Port number of the LDAP service. Typically, this value is 389 for LDAP and 636 for LDAPS. Valid values range from 1 through 65,535.
 - ◆ **Administrator Bind DN** — Distinguished name of the privileged account used to perform operations, such as searching users and groups, on the LDAP directory. Only alphanumeric characters are allowed. Values longer than 60 characters will wrap.
 - ◆ **Bind Password** — Password of the privileged account. Only alphanumeric characters are allowed. Values longer than 15 characters will wrap.
 - ◆ **User Search Path** — Distinguished name of the node at which to begin user searches. Only alphanumeric characters are allowed. Values longer than 40 characters will wrap.
 - ◆ **User Object Class** — Object class identifying users in the LDAP hierarchy. Only alphanumeric characters are allowed. Values longer than 15 characters will wrap.
 - ◆ **User ID Attribute** — Attribute identifying the user login ID within the user object. Only alphanumeric characters are allowed. Values longer than 15 characters will wrap.
 - ◆ **Group Search Path** — Distinguished name of the node at which to begin group searches. Only alphanumeric characters are allowed. Values longer than 40 characters will wrap.
 - ◆ **Group Object Class** — Object class identifying groups in the LDAP hierarchy. Only alphanumeric characters are allowed. Values longer than 15 characters will wrap.
 - ◆ **Group Name Attribute** — Attribute identifying the group name. Only alphanumeric characters are allowed. Values longer than 15 characters will wrap.
 - ◆ **Group Member Attribute** — Attribute indicating group membership for a user within the group object. Only alphanumeric characters are allowed. Values longer than 15 characters will wrap.
 - ◆ **Limit Authentication to members of LDAP group** — Specifies to limit authentication to only members of a specified group.
 - **Group Names** — Name of the authenticated LDAP group. Only alphanumeric characters are allowed. Values longer than 30 characters will wrap.
 - b. To limit authentication to only members of specific groups, click **Show Advanced**, select the option, and then type the **Group Names**, separated by commas.
 - c. Click **Next**.

10. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
11. Click **Finish**.

Viewing authentication authorities

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Authentication** to open the **Authentication** list view.

Use the **Authentication** list view to view and manage authentication settings.

The following properties display:

Authentication Authority States

- ◆ **Local Directory** — When enabled, users can log in as a user from the CST local directory.
- ◆ **Windows OS/AD** — When enabled, users can log in as a user from the Windows local host and/or from the Active Directory domain. This property only displays for Windows installations.
- ◆ **LDAP-SSL** — When enabled, users can log in as a user from the configured LDAP directory.

Windows OS/AD Details (Only displays if Windows OS/AD is enabled, and the advanced **Limit authentication** option is enabled.)

- ◆ **Limit Authentication to Group** — Name of authenticated local Windows or Windows domain group.

LDAP Details (Only displays if LDAP is enabled.)

- ◆ **Server** — Hostname of IP address of the LDAP server used for authentication.
- ◆ **Port** — Port number of the LDAP service. Typically, this value is 389 for LDAP and 636 for LDAPS.
- ◆ **Bind DN** — Distinguished name (DN) of the privileged account used to perform operations, such as searching users and groups, on the LDAP directory.
- ◆ **User Search Path** — Distinguished name of the node at which to begin user searches.
- ◆ **User Object Class** — Object class identifying users in the LDAP hierarchy.
- ◆ **User ID Attribute** — Attribute identifying the user login ID within the object.
- ◆ **Group Search Path** — Distinguished name of the node at which to begin group searches.
- ◆ **Group Object Class** — Object class identifying groups in the LDAP hierarchy.
- ◆ **Group Name Attribute** — Attribute identifying the group name.
- ◆ **Group Member Attribute** — Attribute indicating group membership for a user within the group object.
- ◆ **Limit Authentication to Group** — Name of authenticated LDAP group.

The following control is available:

- ◆ **Edit** — See [Configuring authentication authorities](#) on page 47.

Authorization Rules

Creating authorization rules

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Authorization Rules** to open the **Authorization Rules** list view.
5. Click **Create** to open the **Create Authorization Rules** dialog box.
6. Type a user or group **Name**. User/group names can only contain alphanumeric characters.
7. *Optional:* Select an authentication **Authority**. Possible values are:
 - ◆ **Local Directory** — Specifies to authenticate the user against the Local Authority repository.
 - ◆ **Windows OS** — Specifies to authenticate the user against the Windows local host and.
 - ◆ **Windows AD** — Specifies to authenticate the user against the Active Directory domain.
 - ◆ **LDAP-SSL** — Specifies to authenticate the user against an LDAP directory.
8. *Optional:* Type the authentication **Domain** name. Possible values are based on the authentication authority:

Authority	Domain name
Local directory Windows OS	Unisphere server hostname
Windows AD	Unisphere server domain
LDAP-SSL	LDAP server domain

9. Select a **Role** for at least one Symmetrix system. Note that only Symmetrix systems visible to you are shown on this page.
10. Click **OK**.

Editing authorization rules

This procedure explains how to modify an authorization rule on a Symmetrix system.

Before you begin:

To perform this operation, you must be the Initial Setup User (set during installation), an Administrator, or SecurityAdmin.

To edit authorization rules:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Authorization Rules** to open the **Authorization Rules** list view.
5. Select a rule and click **Edit** to open the **Edit Authorization Rule** dialog box.
6. Make your changes.
7. Click **OK**.

Deleting authorization rules

This procedure explains how to delete authorization rules.

Before you begin:

To perform this operation, you must be the Initial Setup User (set during installation), an Administrator, or SecurityAdmin.

To delete authorization rules:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Authorization Rules** to open the **Authorization Rules** list view.
5. Select the rule and click **Delete**.
6. Click **OK** on the confirmation message.

Viewing authorization rules

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Authorization Rules** to open the **Authorization Rules** list view.

Use the **Authorization Rules** list view to view and manage authorization rules.

The following properties display:

- ◆ **Type** — Whether the user is an individual **User** or **Group**. If a user or group is assigned roles on multiple Symmetrix systems, click the arrow to expand the list.
- ◆ **Name** — User or group name.
- ◆ **Authority** — Authentication authority. Possible values are:
 - **Local Directory** — Directory of users and encrypted passwords stored in a CST .xml file (users only, no groups).
 - **Windows OS** — Local Windows users and groups.
 - **Windows AD** — Windows Active Directory users and groups that are accessed through the SMAS server's domain.
 - **LDAP-SSL** — Users and groups on LDAP server that have been configured the Configure Authorization wizard.
- ◆ **Authentication Domain** — Domain name. Possible values are based on the authentication authority:

Authority	Domain name
Local directory Windows OS	Unisphere server hostname
Windows AD	Unisphere server domain
LDAP-SSL	LDAP server domain
Virtualization domain	Virtualization domain
Any authority	Any

- ◆ **Symmetrix** — Symmetrix system on which the role is assigned.
- ◆ **Role** — Role to which the user or group is mapped. Possible values are:
 - **None** — Provides no permissions.
 - **Monitor** — Performs read-only (passive) operations on a Symmetrix system, excluding the ability to read the audit log or Access Control definitions.
 - **Storage Admin** — Performs all management (active or control) operations on a Symmetrix system and modifies GNS group definitions in addition to all Monitor operations.

- **Admin** — Performs all operations on a Symmetrix system, including security operations, in addition to all StorageAdmin and Monitor operations.
- **Security Admin** — Performs security operations on a Symmetrix system, in addition to all Monitor operations.
- **Auditor** — Grants the ability to view, but not modify, security settings for a Symmetrix system, (including reading the symacl list, and symauth) in addition to all Monitor operations. This is the minimum role required to view the Symmetrix audit log.

The following controls are available:

- ◆ **Create** — See [Creating local users](#) on page 56.
- ◆ **Edit** — See [Editing authorization rules](#) on page 52.
- ◆ **Delete** — See [Deleting local users](#) on page 57.

Local Users

Understanding user authorization

User authorization is a tool for restricting the management operations users can perform on a Symmetrix system. By default, user authorization is enabled for Unisphere users, regardless of whether it is enabled on the Symmetrix system.

When configuring user authorization, an Administrator or SecurityAdmin maps individual users or groups of users to specific roles on Symmetrix systems, which determine the operations the users can perform. These user/role/Symmetrix system mappings (known as authorization rules) are maintained in the symauth users list file, which is located on either a host or Symmetrix system, depending on the Symmetrix Engineuity level.

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 Symmetrix systems from the Unisphere console.

Roles

The following introduces the available roles. For a more detailed look at the permissions that go along with each role, see [Roles and associated permissions](#) on page 60.

- ◆ **None** — Provides no permissions.
- ◆ **Monitor** — Performs read-only (passive) operations on a Symmetrix system.
- ◆ **StorageAdmin** — Performs all management (active or control) operations on a Symmetrix system and modifies GNS group definitions in addition to all Monitor operations.
- ◆ **Administrator** — Performs all operations on a Symmetrix system, including security operations, in addition to all StorageAdmin and Monitor operations.
- ◆ **SecurityAdmin** — Performs security operations on a Symmetrix system, in addition to all Monitor operations.
- ◆ **Auditor** — Grants the ability to view, but not modify, security settings for a Symmetrix system, in addition to all Monitor operations.

A user cannot change their own role so as to remove Administrator or SecurityAdmin privileges from themselves.

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. For more information, see [Operating as the initial setup user](#) on page 20.

Individual and group roles

Users gain access to a Symmetrix 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.

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 (e.g., `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.

Creating local users

This procedure explains how to create local users. Local users have accounts stored locally on the Unisphere server host. (Local users log in to Unisphere with your Username and Password.

Before you begin:

To perform this operation, you must be the Initial Setup User (set during installation), an Administrator, or SecurityAdmin.

To create local directory users:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Local Users** to open the **Local Users** list view.
5. Click **Create** to open the **Create Local Directory** dialog box.
6. Type a user or group name (**User Name**). User/group names are case-sensitive and can only contain alphanumeric characters.
7. Type and confirm a user **Password**. Passwords cannot exceed 16 characters. Only alphanumeric and the following special characters are allowed: `~ ! @ # $ % ^ & * () { } [] ` " < > ' , . ; / ? = + \ |`
8. For each applicable Symmetrix system, assign a role to the user.
9. Click **OK**.

Changing local user passwords

This procedure explains how to change a local user passwords.

Before you begin:

To perform this operation, you must be the Administrator or SecurityAdmin.

To change local directory user passwords:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Local Users** to open the **Local Users** list view.
5. Click **Change Password** to open the **Change Password** dialog box.
6. Type the user's **Current Password**.
7. Type a **New Password** and **Confirm New Password**.
8. Click **OK**.

Deleting local users

This procedure explains how to delete local users.

Before you begin:

To perform this operation, you must be the Initial Setup User (set during installation), an Administrator, or SecurityAdmin.

To delete a user:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Local Users** to open the **Local Users** list view.
5. Select the user and click **Delete** to open the Remove Authorization Rule dialog box.
6. *Optional:* By default, deleting local users removes their associated authorization rules. To override this default behavior, clear the **Remove associated authorization rules** option.
7. Click **OK**.

Adding authorization rules to local users

This procedure explains how to add authorization rules to existing local users.

Before you begin:

- ◆ To perform this operation, you must be the Initial Setup User (set during installation), an Administrator, or SecurityAdmin.
- ◆ Review [Roles and associated permissions](#) on page 60 for information on the permissions that go along with each role.

To add authorizations to local users:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Local Users** to open the **Local Users** list view.
5. Select a user click **View Details** to open the user **Details** view.
6. In the **Related Object** panel, click **Authorization Rules** to open the **Authorization Rules** list view.
7. Click **Add** to open the **Add Authorization** dialog box.
8. Select a **Role** for the user on each applicable Symmetrix system.
9. Click **OK**.

Viewing local users

This procedure explains how to view users with a local Unisphere account.

This procedure explains how to view details on a local directory user:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Local Users** to open the **Local Users** list view.

The **Local Users** list view to allows you to view and manage local users.

The following properties display:

- ◆ **User Name** — User or group name.
- ◆ **Description** — Optional description defined during creation.

The following controls are available:

- ◆ **Create** — See [Creating local users](#) on page 56.
- ◆ **Delete** — See [Deleting local users](#) on page 57.
- ◆ **Change Password** — [Changing local user passwords](#) on page 56
- ◆ **View Details** — [Viewing local users details](#) on next page

Viewing local users details

This procedure explains how to view details on a local user:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Security** to open the **Security** page.
4. Click **Local Users** to open the **Local Users** list view.
5. Select a user and click **View Details** to open the **Local User Details** view.

The **Local User Details** view allows you to view and manage a local user. It contains [Properties](#) and [Related Object](#) panels.

Properties panel:

The following properties display:

- ◆ **User Name** — User or group name.
- ◆ **Description** — Optional description defined during creation. To modify the description, type a new description over the existing and click **Apply**.

The following controls are available:

- ◆ **Delete** — See [Deleting local users](#) on page 57.
- ◆ **Change Password** — [Changing local user passwords](#) on page 56
- ◆ **Apply** — Applies changes made to the **Properties** panel.
- ◆ **Cancel** — Cancels changes made to the **Properties** panel.

Related Objects panel:

The **Related Objects** panel provides links to views for objects associated with the user. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking [Authorization Rules - 2](#) will open a view listing the two authorization rules associated with the user.

Roles and associated permissions

The following table details the permissions that go along with each role in Unisphere for VMAX.

Note The Unisphere for VMAX Initial Setup User has all permissions on a Symmetrix system until an Administrator or SecurityAdmin is added to the Symmetrix system.

Permissions	Roles						
	Admin- istrator	Stor- ageAdmin	Mon- itor	Secu- rityAdmin	Audi- tor	No- ne	Perf Mon- itor
Create/de- lete user accounts	Yes	No	No	Yes	No	No	No
Reset user password	Yes	No	No	Yes	No	No	No
Create roles	Yes	Yes	No	Yes (self excluded)	No	No	No
Change own password	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Manage Symmetrix system	Yes	Yes	No	No	No	No	No
Discover Symmetrix systems	Yes	No	No	Yes	No	No	No
Add/show license keys	Yes	Yes	No	No	No	No	No
Set alerts and Optimizer monitoring options	Yes	Yes	No	No	No	No	No
Release Symmetrix locks	Yes	Yes	No	No	No	No	No
Set Access Controls	Yes	Yes	No	No	No	No	No
Set replication and reservation preferences	Yes	Yes	No	No	No	No	No
Access performance data	Yes	Yes	Yes	Yes	Yes	No	Yes
Start data traces	Yes	Yes	Yes	Yes	Yes	No	Yes

Permissions	Roles						
	Admin- istrator	Stor- ageAdmin	Mon- itor	Secu- rityAdmin	Audi- tor	No- ne	Perf Mon- itor
Set performance thresholds/alerts	Yes	Yes	No	No	No	No	Yes
Create and manage performance dashboards	Yes	Yes	Yes	Yes	Yes	No	Yes

Link and launch

Creating link-and-launch client registrations

This procedure explains how to register a ProSphere® console host as a trusted client with the SMAS server. Once registered, ProSphere users can launch EMC® Unisphere for VMAX without logging in.

Before you begin:

- ◆ To perform this operation, you must be an Administrator or SecurityAdmin.
- ◆ ProSphere users planning on launching EMC® Unisphere for VMAX from their respective clients, must have an Administrator or StorageAdmin role assigned to them within EMC® Unisphere for VMAX.

To create link-and-launch client registrations:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Link and Launch** to open the **Link and Launch** list view.
4. Click **Create** to open the **Register Launch Client** dialog box.
5. Type a unique **Client ID**. Client IDs can be up to 75 alphanumeric characters.
6. Type the **Password** associated with the client ID. Passwords can be up to 75 alphanumeric characters.
7. Retype the password to confirm it.
8. Click **OK**.

Editing link-and-launch client registrations

This procedure explains how to change the password associated with a registered client.

Before you begin:

To perform this operation, you must be an Administrator or SecurityAdmin.

To edit link-and- launch client registrations:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Link and Launch** to open the **Link and Launch** list view.
4. Select a registration, and click **Edit** to open the **Edit Launch Client** Registration dialog box.
5. Type the **Current Password**.
6. Type the **New Password**. Passwords can be up to 75 alphanumeric characters.
7. Retype the new password to confirm it.
8. Click **OK**.

Deleting link-and-launch client registrations

This procedure explains how to delete a client registration from the SMAS server. Once deleted, the client is no longer be able to launch EMC® Unisphere for VMAX.

Before you begin:

To perform this operation, you must be an Administrator or SecurityAdmin.

To delete link-and-launch client registrations:

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Link and Launch** to open the **Link and Launch** list view.
4. Select a registration, and click **Delete**.
5. Click **Yes** in the confirmation box.

Viewing link and launch client registrations

1. From the system selector, select **All Symmetrix**.
2. In the **Home** section's **Common Tasks** panel, click **Administration** to open the **Administration** page.
3. Click **Link and Launch** to open the **Link and Launch** list view.

The **Link and Launch** list view allows you to view and manage link and launch client registrations.

The following property displays:

- ◆ **Client ID** — Unique client ID

The following controls are available:

- ◆ **Create** — [Creating link-and-launch client registrations](#) on page 63
- ◆ **Edit** — [Editing link-and-launch client registrations](#) on previous page
- ◆ **Delete** — [Deleting link-and-launch client registrations](#) on previous page

CHAPTER 3

System Management

This chapter contains the following:

- Monitoring Symmetrix systems 68
- Setting system attributes 70
- Setting CPU I/O resource distribution 72
- Audit log 73
- Hardware management 76
- Job management 97
- Electronic licenses 106
- Symmetrix access controls 118
- Dynamic Cache Partitioning 130

Monitoring Symmetrix systems

The **System Dashboard** provides you with a single place from which to monitor a Symmetrix system, including capacity, hardware, and alert information.

To access the System Dashboard:

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.

System Dashboard:

The **System Dashboard** includes the following components:

System details

- ◆ Symmetrix serial number. To view detailed information for the Symmetrix system, click View Details.
- ◆ **Model** — Symmetrix model number.
- ◆ **Enginuity** — Enginuity version running on the Symmetrix system. When followed by a warning icon (⚠), the Enginuity version is below the minimum suggested version. In which case, you should contact your service provider to schedule an upgrade.
- ◆ **Connection** — Whether the Symmetrix system is local or remote.
- ◆ **Cache Size** — Amount of cache configured on the Symmetrix system.
- ◆ **Health and Maintenance** — Status of the last health check run on the Symmetrix system. Click the link to open the **Health and Maintenance** page.

Capacity panel

Consists of two sections:




- ◆ **Capacity dashboard** — Displays the **Physical Capacity** of all the thin pools on the Symmetrix system and the host-perceived capacity of all thin volumes (**Virtual Capacity**).
- ◆ **Thin Pools list** — Lists the free (green) and allocated (blue) capacity of each thin pool on the Symmetrix system.



Hardware panel

Provides quick links to some of the more important hardware components of the Symmetrix system.

Alerts panel

Consists of two sections:

- ◆ **Alert dashboard** — Displays all the alerts currently on the Symmetrix system, by severity:
 -  Fatal
 -  Critical
 -  Warning

-  Information
-  Normal
- ◆ **Alert list** — Lists the most recent Critical and Fatal alerts on the Symmetrix system and their properties, including:
 - **Severity** — Alert's severity. Possible values are:
 - (1) Fatal
 - (2) Critical
 - (3) Warning
 - (4) Information
 - (5) Normal
 - **Type** — Type of alert. Possible values are Array, Performance, and System.
 - **Description** — Description of the alert.
 - **Time Stamp** — Date/time the alert was created.

Setting system attributes

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Attributes** to open the **Symmetrix Attributes** page.
3. Set or clear any number of the following general attributes:
 - ◆ **Hot Swap Policy** — Specify whether to use global sparing:
 - **Permanent** — Specifies that the spare drive become the active (permanent) drive. In this case, the data is not moved back to the drive that experienced the failure.
 - **Temporary** — Specifies that once the failed drive is swapped, the spare copies all of the data back to the original drive and becomes spare again. This setting only applies to Symmetrix systems running Engenuity versions lower than 5874.
 - ◆ **VCMDB Access Restricted** — Specify whether the VCM database can be accessed on this Symmetrix system.
Select to enable, or clear to disable.
4. Specify whether to use auto-meta feature. Select to enable, or clear to disable.
If Auto Meta is enabled, set the following properties:
 - ◆ **Minimum Meta Capacity** — Type the minimum volume size that will trigger the creation of a meta volume.
For Engenuity 5874 or higher, this value must be less than or equal to 525336 cylinders, if running in 32K compatibility mode; or 262669 cylinders, if running in native mode.
 - ◆ **Member Capacity** — Type the size of the meta members to use when creating meta volumes.
For Engenuity 5874, this value must be less than or equal to 525336 cylinders, if running in 32K compatibility mode; or 262668 cylinders, if running in native mode.
 - ◆ **Configuration** — Type the meta configuration as either **Concatenated** or **Striped** when creating meta volumes.
When enabled and attempting to create a volume larger than the value specified in the **Minimum Meta Capacity** field, or larger than 59 GB, it automatically triggers the creation of a meta volume according to the values specified in the **Member Capacity** and **Configuration** fields.
5. Specify values for the following PAV attributes:
This panel only displays only for Symmetrix systems with either ESCON or FICON directors.
Set the following parameters (z/OS only):
 - ◆ **Alias Limit** — Type the maximum number of aliases that can be assigned to a volume.
 - ◆ **Mode** — Select the one of the following PAV types:
 - **NoPAV** — PAV not configured for the Symmetrix system.
 - **Standard PAV** — Volumes with static aliasing.
 - **Dynamic Standard PAV** — Standard PAV volumes with dynamic aliasing.

6. Specify values for the following SRDF/A attributes:
 - ◆ **Maximum Cache Usage (Percent)** — Type the maximum percentage of system write-pending cache slots for all RDF/A sessions. Valid values are 0 to 100.
 - ◆ **Maximum Host Throttle (Secs)** — Type the maximum percentage of system write-pending cache slots for all RDF/A sessions. Valid values are 0 to 65535.
7. Specify values for the following SRDF director attributes:

These attribute settings control mixed I/O workloads on the same SRDF director.

 - ◆ **SRDF Director CPU Resource Distribution** — Whether the director CPU resource distribution feature is enabled or disabled.
 - ◆ **Synchronous I/O Percent** — Type the percentage of director resources to allocate for SRDF synchronous I/Os.
 - ◆ **Asynchronous I/O** — Type the percentage of director resources to allocate for SRDF asynchronous I/Os.
 - ◆ **Copy I/O Percent** — Type the percentage of director resources to allocate for copy I/Os.

Percentages for synchronous, asynchronous, and copy I/O must add up to 100. Percentage values can be set even if the **SRDF Director CPU Resource Distribution** is disabled.
8. Specify values for the following DCP values:
 - ◆ **Cache Partition Status** — Whether the dynamic cache partition feature is enabled, disabled, or in analyze mode.

Analyze mode is a tool for determining the amount of cache your applications are consuming, prior to enabling the cache partitioning feature. For more on analyze mode, see [Running in analyze mode](#) on page 131.
 - ◆ **Empty Partition Status** — Whether to preserve or automatically remove empty cache partitions. By default, all empty partitions are removed if there is no partition configuration activity for 4 hours.
9. Click **Apply**.

Setting CPU I/O resource distribution

The following explains how to set parameters to control mixed I/O workloads on an SRDF director. Note that values specified here will override the default values defined in See [Setting system attributes](#) on page 70.

1. Select the Symmetrix system.
2. Select **System** > **Dashboard**.
3. In the **Hardware** panel click **RDF Directors** to open the **RDF Directors** list view.
4. Select an RDF director, and click **CPU I/O Settings** to open the **Set SRDF Director CPU I/O** dialog box.
5. Type values for any number of the following parameters. Note that the percentages must add up to 100.
 - ◆ **Synchronous I/O Percent** — Type the percentage of director resources to allocate for synchronous I/Os.
 - ◆ **Asynchronous I/O** — Type the percentage of director resources to allocate for asynchronous I/Os.
 - ◆ **Copy I/O Percent** — Type the percentage of director resources to allocate for copy I/Os.
6. To reset the parameters to their defaults, click **Reset**.
7. Click **OK**.

Audit log

Viewing the Symmetrix audit log

The Symmetrix audit records come from the SYMAPI database and include all actions taken on that Symmetrix system. The audit log resides on the Symmetrix system and currently has a maximum size of 40 MB. Once the 40 MB limit is reached, the log begins to overwrite itself.

Before you begin:

- ◆ Because Symmetrix audit logs can be very large, the list view opens pre-filtered to display the most recent 512 records. For information on filtering the listed records, refer to [Filtering audit log records](#) on next page..
- ◆ The user must have Auditor or higher permissions.

To view the Symmetrix audit log:

1. Select the Symmetrix system.
2. Select **System** › **Audit Log** to open the **Audit Log** list view.

The following properties display:

- ◆ **Record** — Unique identifier for the audit entry.
- ◆ **Date** — Date the audit entry was made.
- ◆ **Application** — Application operating on the Symmetrix system.
- ◆ **Username** — Name of the user operating on the Symmetrix system.
- ◆ **Host** — Host operating on the Symmetrix system.
- ◆ **Function Class** — Generic audit category for the operation on the Symmetrix system.
- ◆ **Action Code** — Specific audit code for the operation on the Symmetrix system.
- ◆ **OS Name** — Operating system running on the host.
- ◆ **Activity ID** — Activity ID for audit record.

The following control is available:

- ◆ **View Details** — [Viewing Symmetrix audit log details](#) below.

Viewing Symmetrix audit log details

To view details on a specific audit log record:


1. Select the Symmetrix system.
2. Select **System > Audit Log** to open the **Audit Log** list view.
3. Select a record and click **View Details** to open its **Details** view.

The following properties display:

- ◆ **Record Number** — Unique identifier for the audit entry.
- ◆ **Text** — Summary of the Symmetrix operation.
- ◆ **Time** — Date the audit entry was made.
- ◆ **Application ID** — Application operating on the Symmetrix system.
- ◆ **Username** — Name of the user operating on the Symmetrix system.
- ◆ **Function Class** — Generic audit category for the operation on the Symmetrix system.
- ◆ **Action Code** — Specific audit code for the operation on the Symmetrix system.
- ◆ **Host** — Host operating on the Symmetrix system.
- ◆ **Records in Seq** — One Symmetrix operation can be represented by a sequence of audit records. This is the total number of records in this particular audit sequence.
- ◆ **Offset in Seq** — Audit entry number within the audit sequence.
- ◆ **Vendor ID** — Vender of the application operating on the Symmetrix system.
- ◆ **Application Version** — Version of the application operating on the Symmetrix system.
- ◆ **API Library** — SYMAPLI library type.
- ◆ **API Version** — SYMAPI version number.
- ◆ **OS Name** — Operating system running on the host.
- ◆ **OS Revision** — Specific revision of the operating system.
- ◆ **Client Host** — Client/Server only.
- ◆ **Activity ID** — Activity ID for audit record.
- ◆ **Process ID** — ID of the process that logged the record.
- ◆ **Task ID** — ID of the task that logged the record.

Filtering audit log records

To filter the logs listed in an audit log list view:

1. Click the filter icon  to open the **Auto Log Filter** dialog box.
2. Select or type a value for any number of the following criteria, and then click **OK**.

Range

- ◆ **By Date/By Record Number**— Specifies whether to filter the records by date/time or by

record number and enables the appropriate fields:

- **Start Date** — Specifies the start date and time to use in the range. Click the calendar icon to pick the date. Use arrow controls to select the time. Select **Earliest** to set the Start Date to the date/time the first record was logged.
 - **End Date** — Specifies the end date and time to use in the range. Click the calendar icon to pick the date. Use arrow controls to select the time. Select **Latest** to set the End Date to the date/time the last record was logged.
 - **Start Record** — Specifies the first record ID to use in the range.
 - **End Record** — Specifies the last record ID to use in the range.
- ◆ **Advanced Filtering:**
- **Function Class** — Filters the list for records with the specified audit category.
 - **Action Code** — Filters the list for records with the specified action code.
 - **Activity ID** — Filters the list for records with the specified activity ID.
 - **Host** — Filters the list for records logged by a specific host.
 - **Vender ID** — Filters the list for records logged by a specific vendor.
 - **Application ID** — Filters the list for records logged by a specific application.
 - **User Name** — Filter the list for records logged by a specific user.
 - **Affected Volumes** — Filter the list for records logged by specific volumes.
1. To clear the filter, open the **Audit Log Filter** dialog box, click **Clear All**, and then **OK**.

Hardware management

Converting directors

This procedure explains how to convert directors (FA to RF and from RF to FA).

1. Select a Symmetrix system.
2. Select **System > Dashboard**.
3. Do the following, depending on whether you are converting front end directors or RDF directors:

Front end directors:

- a. To convert an *FA* director, click **Front End Directors** to open the **Front End Directors** list view.
- b. Select a director, and click **Convert FA to RF** to open the **Convert Front End Director to RDF Director** dialog box.
- c. Click either of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to convert the director now.

RDF directors:

- a. To convert an *RDF* director, click **RDF Directors** to open the **RDF Directors** list view.
 - b. Select a director and click **Convert RF to FA** to open the **Convert RDF Director to Front End Director** dialog box.
 - c. Click either of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to convert the director now.
1. *Optional:* Monitor the Alerts list view (**System > Alerts**) to monitor the conversion process. The following alerts are typical during a conversion:
 - ◆ Director status has changed to Failed.
 - ◆ Port state has changed to Offline.
 - ◆ Port state has changed to Offline.
 - ◆ Director state has changed to Online
 - ◆ Port state has changed to Not Present.
 - ◆ Port state has changed to Online.

Setting director port attributes

1. Select the Symmetrix system.
2. Select **System** > **Dashboard**.
3. From the **Hardware** panel, click **Front End Directors** to open the **Front End Directors** list view.
4. Select the director, and click **Set Port Attributes** to open the **Set Port Attributes** dialog box.
5. *Optional:* Select a port whose flag settings you want to copy.
6. Select/clear any number of the following attributes:

Note The following lists the attributes for all port types. Therefore, depending on the port type, some of the attributes may not appear in the dialog box.

ACLX	Enables the port to be added to a port group.
Common Serial Number	Enables multi-path configurations or hosts that need a unique serial number to determine which paths lead to the same volume.
Unique WWN	Ensures unique World Wide Names (WWNs) within the fiber environment (uses Symmetrix serial numbers and port numbers). This is enabled by default for all environment configuration changes and new environments. When disabled, you don't have to change WWNs.
Init Point to Point	Specifies a point-to-point (direct or switched) topology in the initialization sequence. When disabled (default), it is initialized as an arbitrated loop.
Volume Set Addressing	Enables the volume set addressing mode. When using volume set addressing, you must specify a 4-digit address in the following range: (0)000-(0)007, (0)010-(0)017,... to a maximum of (0)FF0-(0)FF7 The first digit must always be set to 0 (Symmetrix systems do not currently support the upper range of volume set addressing), the second digit is the VBus number, the third digit is the target, and the fourth digit is the LUN.
Avoid Reset Broadcast	Enables a SCSI bus reset to only occur to the port that received the reset (not broadcast to all channels).
Negotiate Reset	When enabled for AS/400 hosts, this flag forces a SCSI negotiation by the Symmetrix system after a SCSI reset, an error, or a bus volume reset.
Enable Auto Negotiate	Allows two fibre ports to handshake and settle on an optimal speed for data transfer.
Environ Set	Enables the environmental error reporting by the Symmetrix system to the host on the specific port.

Disable Q Reset on UA	When enabled, a Unit Attention (UA) that is propagated from another director does not flush the queue for this volume on this director. Used for hosts that do not expect the queue to be flushed on a 0629 sense (only on a hard reset).
Soft Reset	Supports SCSI soft reset on a Symmetrix port when enabled for a Bull/GCOS-7 host.
SCSI 3	Alters the inquiry data (when returned by any volume on the port) to report that the Symmetrix system supports the SCSI-3 protocol. When disabled, the SCSI-2 protocol is supported.
SCSI Support1 (OS2007)	Provides a stricter compliance with SCSI standards for managing volume identifiers, multi-port targets, unit attention reports, and the absence of a volume at LUN 0.
SPC2 Protocol Version	This flag should be enabled (default) in a Windows 2003 environment running Microsoft HCT test version 12.1. When setting this flag, the port must be offline.
HP3000 Mode	Causes the Symmetrix port to return a SCSI busy state, instead of a 0B44 sense code when a xx3C error occurs. (Applies to HP MPE 5.0 and Enginuity 5062 and lower).
Sunapee	Enables the Sunapee option on the port, for SUN PDB clusters.
Siemens	Returns in the sense data error 0B48, instead of 0B44 for normal behavior. (Applies to Siemens R-series platforms only).
Sequent	Sets the task timeout window to be 15 seconds before aborting a process, and the Symmetrix system on the specified port does not force wide or synchronous negotiations. Also, a busy status returns, instead of a 0B44h when aborting a command on a timeout.
Server on AS400	Indicates the port is to behave as a server returning server inquiry data, rather than AS/400 data. (Applies to AS/400 platforms only).
Enable AS400	Indicates whether AS/400 is enabled.
Open VMS	Enables an Open VMS fiber connection.
iSCSI IPv4 Address	Identifies the port's iSCSI IP address (IPv4).
iSCSI IPv4 Default Gateway	Identifies the port's default gateway address (iSCSI).
iSCSI IPv4 Netmask	Identifies the port's netmask address (iSCSI).
iSCSI Initiator Session	Identifies the port's initiator session ID (iSCSI).

iSCSI IPv6 Address	Identifies the port's iSCSI IP address (IPv6).
iSCSI IPv6 Net Prefix	Identifies the port's iSCSI net prefix (IPv6). Possible values range from 0 to 127. The default value is 64.

- Click **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.

Enabling and disabling director ports

- Select the Symmetrix system.
- Select **System > Dashboard**.
- From the **Hardware** panel, click **Front End Directors** to open the **Front End Directors** list view.
- Select one or more ports, and click **Enable** to enable them, or **Disable** to disable them.
- Click **OK** in the confirmation message.

Performing system health checks

This procedure explains how to inspect the general state of a Symmetrix system.

Health check performs the following tests:

- ◆ **Vault State Test** — Verifies the ability of the system to save data in case of a power failure.
- ◆ **Spare Drive Test** — Verifies that spare drives are available in case of a drive failure.
- ◆ **Memory Test** — Verifies that the memory is reporting no errors or disabled banks.
- ◆ **Locks Test** — Verifies that there are no software locks present.
- ◆ **Emulations Test** — Verifies that all directors are loaded with the same Enginuity release as that on the service processor.
- ◆ **RDF Test** — Verifies that all SRDF links are online.
- ◆ **Environmental Test** — Verifies that internal environmental components (power supplies, fans, batteries, etc.) are reporting no errors.
- ◆ **Battery Test** — Verifies that the most-recent battery test reported no errors.
- ◆ **General Tests** — Checks for any abnormal conditions in the following areas: volume status, director status, hung upgrade, code table integrity, directors running same code.

Before you begin:

- ◆ To perform this procedure you must be an Administrator or Storage Admin.
- ◆ The Symmetrix system must be running Enginuity 5875 or higher.

To perform a health check on a Symmetrix system:

1. Select a Symmetrix system
2. Select **System** > **Health and Maintenance** to open the system **Health and Maintenance** page.
3. Click **Health Check** .
The **Health Check** view opens displaying results of previously run health checks, if any.
4. Click **Run**.
5. Determine the health of the Symmetrix system by examining the **Status** and **Results** fields.

Replacing failed drives

This procedure explains how to locate and replace failed Symmetrix disk drives.

The following are the high-level steps for completing this procedure:

[Step 1: Locate the failed drive](#) below

[Step 2: Replace the failed drive](#) on the facing page

[Step 3: Verify the drive replacement](#) on page 83

Before you begin:

- ◆ To perform this procedure you must be an Administrator or Storage Admin.
- ◆ The Symmetrix VMAX Series system with the failed disk drive must be locally attached to the SMAS server, running Enginuity 5876 Q2 2013 SR or higher, and have the CRU flag enabled.
- ◆ Ensure that there are no other disk replacements occurring on the Symmetrix system..
- ◆ Ensure that you have the replacement drive and Symmetrix door key.
- ◆ **Once you run the disk replacement job, you must replace the failed drive within 30 minutes to complete the disk replacement process. If you do not replace the drive within this time frame, the job will end, and you will need to restart the process.**

Step 1: Locate the failed drive

1. Select a Symmetrix system
2. Select **System** > **Health and Maintenance** to open the system **Health and Maintenance** page.
3. Click **Spare Drive Replacement** to open the **Spare Drive Replacement** list view.
4. Select the drive and click **Replace** to open **Spare Drive Replacement** dialog box.
5. Do one of the following:
 - ◆ Click **Run Now** to replace the drive now.
 - ◆ Expand **Run Now** and click **Add to Job List** to replace the drive at a later time, as described in [Running jobs](#) on page 100 [Scheduling jobs](#) on page 100.

Regardless of the method you use to run the task, you have 30 minutes to locate and replace the failed drive. If you do not replace the drive within this time frame, the job will end, and you will need to restart the process.

Once you run the job, the **Spare Drive Replacement Instructions** dialog will open, displaying instructions and a map for locating the drive.

6. **Print** or **Export** the dialog contents, and then click **Close**.
7. Carefully following the spare drive replacement instructions you just printed/exported to locate the drive.

Step 2: Replace the failed drive


It is recommended that before you begin any of the following procedures that you attach the clip of an ESD wristband to bare metal on the storage bay, and secure the wristband around your wrist with the metal button against your skin.

To replace a failed disk drive:

1. Access the disk drive. Do the following, depending on the Symmetrix model:

Symmetrix VMAX Series system:

- a. Open the front door of the Symmetrix system and press the two spring latches on the DAE cover. Pull the cover outward and remove.
- b. Place the DAE cover aside for reinstallation later.




Symmetrix VMAX SE system:


Open the front door of the Symmetrix system.

Symmetrix VMAX 10K system:


Remove EMI front panel by pressing the side tabs and pulling straight forward.



1. Remove the disk drive:
 - a. Release the latch on front of the disk you want to remove.
 - b. Gently pull the disk out 1 inch and wait 30 seconds for the disk to stop spinning.
 - c. Remove the disk and place it on an anti-static surface.



2. Install the disk drive:
 - a. Align the drive with the guides in the slot.
 - b. With the disk drive latch fully open, gently push the drive into the slot.
The disk drive latch will begin to rotate downward when its tabs meet the enclosure chassis.
 - c. Push the handle down to engage the latch. After the latch is engaged, push firmly on the bottom of the drive to verify that the disk is properly seated.



3. Reinstall the DAE cover or EMI front panel, or close the door, depending on the Symmetrix model:

Symmetrix VMAX Series system:


- a. Depress the two spring latches on the DAE cover and align them with the slots on the chassis.
- b. Release the tabs to secure the cover. Close the front door of the storage bay.

Symmetrix VMAX SE system:

Close the door.

Symmetrix VMAX 10K system:

At the front of the cabinet, install the EMI front panel by pressing the side tabs and snapping into place.



Step 3: Verify the drive replacement

Once you have completed the disk drive replacement, the lights will return to normal after several seconds. At this point you should return to the Unisphere for VMAX console to ensure that the disk replacement completed without issues.

To verify that you successfully replaced the drive, look for the following message in the job list (for instructions on viewing the job list, refer [Viewing the job list](#) on page 102):

```
Spare Drive Replacement SUCCEEDED
```

If any other messages appear, follow the instructions.

Locating Symmetrix VMAX systems

The following procedures explain how to locate a Symmetrix VMAX system by flashing/blinking its bay door lights.

With this feature you can either control the LEDs on all the bay doors or on individual bay doors.

Before you begin:

You cannot:

- ◆ Set the system bay LEDs to **Blue Flash** if any bay led state is anything other than **On**.
- ◆ Reset all LEDs when any LED state is **N/A** or **Blink Fast**.
- ◆ Set a bay LED to **Slow Blink** if any other bay LED is **Off**, **Blink Slow**, **Fast Blink**, or **Blue Flash**.
- ◆ Set a bay LED to **On** if the current LED state is anything other than **Off**, **Blink Slow**, or **Blue Flash**.

To control all the bay LEDs:

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.
3. In the **System Component** panel, right-click the Symmetrix system and select either of the following:
 - ◆ To set all of the system bay LEDs to Blue Flash, select **Blue Flash On**.
 - ◆ To reset all of the bay LEDs to the solid (On) state, select **Reset all LEDs**.

To control individual bay LEDs:

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.
3. In the **System Component** panel, right-click the component and select either of the following:
 - ◆ To blink the bay LED, select **Blink On**.
 - ◆ To reset the bay LED to the solid (On) state, select **Blink Off**.

Viewing hardware components

Viewing Symmetrix hardware components

The **Symmetrix Hardware** dashboard, allows you to view a system configuration and sub-components.

From this view you can:

- ◆ View system sub-components in a tree view or details view.
- ◆ Locate system bay sub-components, such as engines, power supplies, and drive enclosures.
- ◆ View the current health or state of each system component.
- ◆ Control bay door LED blinking to identify a system bay.

To access the Symmetrix Hardware dashboard:

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.

The dashboard opens with the top-level system configuration as the default view.

3. Drill-down into the **System Components** menu to view details on the system's components.

Viewing overall system configuration

This procedure explains how to view the overall configuration of a Symmetrix system.

To view the overall system configuration:

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.

The dashboard opens with the top level system configuration as the default view.

The tree view lists the system and drive bays, and can be expanded to list sub-components.

The details view, located below the system graphic, displays the system properties.

Viewing drive bay sub-components

This procedure explains how to view the drive bay sub-components of a Symmetrix system.

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.

The dashboard opens with the top-level system configuration as the default view.

3. To view the drive bay sub-components, expand the drive bay in the tree view.
4. The following sub-components are listed and are viewed as follows:

Drive Enclosures

- a. To view drive enclosures in the bay, expand **Drive Enclosures**.

All drive enclosures in the bay are listed as DE-1 - DE-16.

Note Depending on the system configuration, the system bay may also contain drive enclosures, where they are listed as DE-1 - DE-10.

- b. To display the location of an enclosure in the bay, select one from the list and the location is highlighted in the graphic.

The details view, located below the system graphic, displays the drive enclosure properties and status.

Standby Power Supplies

- a. To view the Standby Power Supplies in the bay, expand **Standby Power Supplies**.

All drive enclosures in the bay are listed (SPS-1 - SPS-4).

- b. To display the location of the power supply in the bay, select one from the list and the location is highlighted in the graphic.

The details view, located below the system graphic, displays the power supply properties and status.

Viewing system bay sub-components

This procedure explains how to view Symmetrix system bay sub-components:

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.

The dashboard opens with the top-level system configuration as the default view.

3. To view the system bay sub-components, expand the system bay in the tree view.
4. The following sub-components are listed and are viewed as follows, but may vary depending on the system configuration:

Engines

Lists all the engines (1 - 8) in the bay. Selecting one will display its properties and status in the details view.

Drive Enclosures

Lists all the drive enclosures (DE-1 - DE-10) in the bay. Selecting one will display its properties and status in the details view.

Standby Power Supplies

Lists all the Standby Power Supplies in the bay. Selecting one will display its properties and status in the details view.

Matrix Interface Board Enclosures

Lists all the Matrix Interface Board Enclosures (MIBE-1A MIBE-1B and MIBE-2A MIBE-2B) in the bay. Selecting one will display its properties and status in the details view.

Servers

Lists the server in the bay. Selecting it will display its properties in the details view.

KVM

Lists the KVM in the bay. Selecting it will display its properties in the details view

Uninterruptible Power Supplies

Lists the Uninterruptible Supply in the bay. Selecting it will display its properties in the details view

Viewing system sub-component health status




This procedure explains how to view the health status of Symmetrix components:

1. Select the Symmetrix system.
2. Select **System > Health and Maintenance > Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.

The dashboard opens with the top-level system configuration as the default view.

The health icon located next to each component in the tree view, shows the health state of system sub-components. The health icon state is propagated up through the tree view to the top (system) level.

There are three possible health states that are displayed:

- ◆ **Green**  — OK - component is healthy
- ◆ **Yellow**  — Component health has degraded
- ◆ **Red**  — Component has failed

3. To locate a degraded or failed component, drill down to find the faulty component.

Viewing back-end directors

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **Back End Directors** to open the **Back End Directors** list view.

Use this view to view and manage the back end directors.

The following properties display:

- ◆ **Director Identifier** — Director name
- ◆ **Type** — Director type
- ◆ **Numeric** — Port numeric ID
- ◆ **Slot** — Slot number
- ◆ **Director Serviced Hypers** — Number of hyper volumes serviced by the director
- ◆ **Protocol** — Transmission protocol for the director port
- ◆ **Director Status** — Director status
- ◆ **Port Interface** — Port interface ID
- ◆ **Port Status** — Indicates port status, ON or OFF

The following controls are available:

- ◆ **View Details** — See [Viewing back-end director details](#) on next page.

Viewing back-end director details

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **Back End Directors** to open the **Back End Directors** list view.
4. Select a director from the list and click **View Details** to open the back end director **Details** view.

The backend director **Details** view allows you to view details on a backend director. This view contains [Properties](#), [Related Objects](#), [Performance Views](#) panels.

Properties panel

The following properties display:

- ◆ **Director Identifier** — Director name.
- ◆ **Port Interface** — Port interface ID.
- ◆ **Type** — Director type.
- ◆ **Director Status** — Director status.
- ◆ **Director Serviced Hypers** — Number of hyper volumes serviced by the director.
- ◆ **Port Status** — Indicates port status, ON or OFF.
- ◆ **Port Serviced Hypers** — Number of hyper volumes serviced by the port.

Related Objects panel

The **Related Objects** panel links you to views displaying objects contained in and associated with the director. Each link is followed by a number, indicating the number of objects in the corresponding view.

Performance Views panel

The **Performance** panel links you to the performance monitor and analyze views for the director.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Viewing external directors

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **External Directors** to open the **External Directors** list view.

Use this view to view the external directors.

The following properties display:

- ◆ **Director Identifier** — Director name
- ◆ **Port** — Port ID
- ◆ **Type** — Director type
- ◆ **Numeric** — Port numeric ID
- ◆ **Slot** — Slot number
- ◆ **Director Serviced Hypers** — Number of hyper volumes serviced by the director
- ◆ **Director Status** — Director status
- ◆ **Port Status** — Indicates port status, ON or OFF

The following control is available:

- ◆ **View Details** — [Viewing external director details](#) on next page

Viewing external director details

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **External Directors** to open the **External Directors** list view.
4. Select the director and click **View Details** to open its **Details** view.

The external director **Details** view allows you to view details on an external director. This view contains [Properties](#), [Related Objects](#), [Performance Views](#) panels.

Properties panel

The following properties display: gftg

- ◆ **Director Identifier** — Director name
- ◆ **Port** — Port ID
- ◆ **Type** — Director type
- ◆ **Director Status** — Director status
- ◆ **Port Status** — Indicates port status, ON or OFF
- ◆ **Number of Volumes** — Number of volumes serviced by the director.
- ◆ **WWN** — Port WWN.

Related Objects panel

The **Related Objects** panel links you to views displaying objects associated with the director. Each link is followed by a number, indicating the number of objects in the corresponding view.

Performance Views panel

The **Performance** panel links you to the performance analyze views for the director.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Viewing system front-end directors

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **Front End Directors** to open the **Front End Directors** list view.

Use this view to view and manage the front end directors.

The following properties display:

- ◆ **Director Identifier** — Director name.
- ◆ **Director Status** — Director status.
- ◆ **Port** — Port number.
- ◆ **Port Status** — Indicates port status, ON or OFF.
- ◆ **Type**—Director transmission protocol.
- ◆ **Port ID** — Port ID.
- ◆ **Numeric** — Port numeric ID.
- ◆ **Slot** — Slot number.
- ◆ **Number of Volumes** — Number of volumes mapped to the director port.
- ◆ **VCM State** — Indicates whether VCM device state is enabled or disabled.
- ◆ **ACLX** — Indicates whether access control logic enabled or disabled.

The following controls are available:

- ◆ **View Details** — See [Viewing system front-end director details](#) below.
- ◆ **Enable** — See [Enabling and disabling director ports](#) on page 79.
- ◆ **Disable** — See [Enabling and disabling director ports](#) on page 79.
- ◆ **Set Port Attributes** — See [Setting director port attributes](#) on page 77.
- ◆ **Convert FA to RF** — See [Converting directors](#) on page 76.
- ◆ **Set ORS Ceiling** — See [Setting Open Replicator ceiling](#) on page 419.

Viewing system front-end director details

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **Front End Directors** to open the **Front End Directors** list view.
4. Select a director and click **View Details** to open the front end director **Details** view.

The **Front End Director Details** view allows you to view and manage a front end director. It contains [Properties](#), [Related Objects](#), and [Performance Views](#) panels.

Properties panel

The following properties display:

- ◆ **Director Identifier** — Director name.
- ◆ **Port** — Port number.
- ◆ **Type**—Director transmission protocol.
- ◆ **Director Status** — Director status.
- ◆ **Port Status** — Indicates port status, ON or OFF.
- ◆ **Number of Volumes** — Number of volumes.
- ◆ **Mapped Devices (including meta members)** — number of mapped devices.
- ◆ **Port ID** — Port ID.
- ◆ **Fibre Channel Loop** — Fibre channel loop.
- ◆ **VCM State** — Indicates whether VCM is enabled or disabled.
- ◆ **ACLX** — Indicates whether ACLX is enabled or disabled.
- ◆ **Common Serial Number** — Indicates whether Common Serial Number feature is enabled or disabled.
- ◆ **Unique WWN** — Indicates whether Unique WWN feature is enabled or disabled.
- ◆ **Init Point to Point**—Indicates whether Init Point to Point is enabled or disabled.
- ◆ **Volume Set Addressing** — Indicates whether Volume Vet Addressing is enabled or disabled.
- ◆ **Avoid Reset Broadcasting** — Indicates whether Avoid Reset Broadcasting feature is enabled or disabled.
- ◆ **Negotiate Reset** — Indicates whether Negotiate Reset feature is enabled or disabled. This feature is used for AS/400 systems only (default is off).
- ◆ **Enable Auto Negotiate** — Indicates whether Enable Auto Negotiate feature is enabled or disabled.
- ◆ **Environ Set** — Indicates whether environmental error reporting feature is enabled or disabled.
- ◆ **Disable Q Reset on UA** — Indicates whether the Disable Q Reset on UA (Unit Attention) is enabled or disabled.
- ◆ **Soft Reset** — Indicates whether Soft Reset feature is enabled or disabled for a Bull/GCOS-7 host.
- ◆ **SCSI 3** — Indicates whether SCSI-3 protocol is enabled or disabled. When disabled, the SCSI 2 protocol is supported.
- ◆ **SCSI Support1(OS2007)** — Indicates whether SCSI Support1 is enabled or disabled.
- ◆ **Non Participate** — Indicates whether Non Participate feature is enabled or disabled.

- ◆ **SPC Protocol Version** — Indicates whether SPC Protocol Version feature is enabled or disabled. This feature is used for Windows 2003 environment running Microsoft HCT test version 12.1.
- ◆ **HP 3000 Mode** — Indicates whether HP 3000 Mode is enabled or disabled. This feature is used for HP MPE 5.0 and Engenuity levels 5062 and earlier.
- ◆ **Sunapee** — Indicates whether Sunapee feature is enabled or disabled. This feature is used for SUN PDB clusters.
- ◆ **Siemens** — Indicates whether Siemens feature is enabled or disabled. This feature is used for Siemens R-series platforms.
- ◆ **Server of AS400** — Indicates whether Server of AS400 feature is enabled or disabled. This feature is used For AS/400 platforms.
- ◆ **Enable AS400** — Indicates whether Enable AS400 is enabled or disabled.
- ◆ **Open VMS** — Indicates whether Open VMS is enabled or disabled.
- ◆ **Recopy Max Ceiling** — Maximum recopy ceiling value.
- ◆ **Recopy Ceiling** — Recopy ceiling value.
- ◆ **Recopy ACU Ceiling** — Recopy ACU ceiling value.
- ◆ **Negotiated Speed (GB/Second)** — Negotiated speed in GB/Second.

The following controls are available:

- ◆ **Enable** — See [Enabling and disabling director ports](#) on page 79.
- ◆ **Disable** — See [Enabling and disabling director ports](#) on page 79.
- ◆ **Set Port Attributes** — See [Setting director port attributes](#) on page 77.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the director. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Mapped Volumes - 2** will open a view listing the two volumes mapped to the director.

Performance Views panel

The **Performance** panel links you to the performance monitor and analyze views for the director.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Viewing RDF directors

1. Select the Symmetrix system.
2. Select **System** > **Dashboard**.
3. In the **Hardware** panel click **RDF Directors** to open the **RDF Directors** list view.

Use this view to view and manage the RDF directors.

The following properties display:

- ◆ **Director Identifier** — Director ID
- ◆ **Type**—Director type
- ◆ **Numeric** — Port numeric ID
- ◆ **Slot** — Slot number
- ◆ **#RDF groups** — Number of RDF groups
- ◆ **Port** — Port ID
- ◆ **Director Status** — Director status
- ◆ **IPSec Policies** — Number of IPSec Policies
- ◆ **Port Status** — Indicates port status, ON or OFF

The following controls are available:

- ◆ **View Details** — See [Viewing RDF director details](#) below.
- ◆ **Convert RF to FA** — See [Converting directors](#) on page 76.
- ◆ **CPU I/O Setting** — See [Setting CPU I/O resource distribution](#) on page 72.

Viewing RDF director details

1. Select the Symmetrix system.
2. Select **System** > **Dashboard**.
3. In the **Hardware** panel, click **RDF Directors** to open the **RDF Directors** list view.
4. Select a director from the list and click **View Details** to open its **Details** view.

The **RDF Director Details** view allows you to view RDF director details. It contains [Properties](#), [Related Objects](#), and [Performance Views](#) panels.

Properties panel

The following properties display:

- ◆ **Director Identifier** — Director name
- ◆ **Type** — Director type
- ◆ **Director Status** — Director status
- ◆ **Port** — Port ID

- ◆ **WWN** — World wide name
- ◆ **RDF (RA) Group Attributes - Farpoint** — Indicates whether or not the communication protocol is Farpoint
- ◆ **Prevent Automatic RDF Link Recovery** — Indicates whether preventing automatic link recovery is enabled or disabled
- ◆ **Prevent RAs Online on Power Up** — Indicates whether preventing RA to be online on power up is enabled or disabled
- ◆ **Protocol Type** — Director transmission protocol type
- ◆ **RDF Software Compression Supported** — Indicates whether RDF Software Compression is supported or not supported.
- ◆ **RDF Software Compression** — Indicates whether RDF Software Compression is enabled or disabled
- ◆ **RDF Hardware Compression Supported** — Indicates whether RDF Hardware Compression is supported or not supported.
- ◆ **RDF Hardware Compression** — Indicates whether RDF Hardware Compression is enabled or disabled
- ◆ **Port Status** — Indicates port status, ON or OFF
- ◆ **IPV4 Address**—IPV4 Address
- ◆ **IPV6 Address** — IPV6 Address
- ◆ **IPV6 Prefix**—IPV6 Prefix
- ◆ **IPV4 Default Gateway** — IPV4 Default Gateway
- ◆ **IPV4 Domain Name** — IPV4 Domain Name
- ◆ **IPV4 Netmask** — IPV4 Netmask
- ◆ **SCSI Initiator Session** — SCSI Initiator Session
- ◆ **Negotiated Speed (GB/Second)** — Negotiated speed in GB/Second
- ◆ **Synchronous I/O Percent** — Percentage of SRDF director CPU resources allocated to synchronous I/Os.
- ◆ **Asynchronous I/O Percent** — Percentage of SRDF director CPU resources allocated to asynchronous I/Os.
- ◆ **Copy I/O Percent** — Percentage of SRDF director CPU resources allocated to copy I/Os.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the director. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **RDF Groups - 2** will open a view listing the two RDF groups in which the director is a member.

Performance Views panel

The **Performance** panel links you to the performance monitor and analyze views for the RDF director.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Job management

Managing job lists

When you perform certain configuration tasks on the Symmetrix 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 named **Add to Job List**.

EMC® Unisphere for VMAX includes a job list view, from which you can view and manage the job list for a Symmetrix system.

The following table contains the configuration tasks that are held for processing in a job list:

Table 1: Job list configuration tasks

Function	Configuration task
System	Set Symmetrix attributes Add Hot Spare Remove Hot Spare
Volumes	Create volumes Duplicate volumes Change volume configuration Set volume attributes Change volume to RDF configuration
Meta volumes	Create/Dissolve meta volumes Convert meta configuration Add/Remove meta members
SAVE volumes/pools	Creating SAVE volumes Add/remove SAVE pool volumes Enable/Disable SAVE volumes
Mapping	Map volumes Unmap volumes
z/OS	Create CKD meta Map volumes Unmap volumes Copy mapping Assign/Remove alias address
Port/HBA	Set port flags

Making configuration changes safely

Before making configuration changes, you must thoroughly understand your Symmetrix configuration. The following guidelines establish safe disciplines as you begin any change that can impact stored data:

- ◆ Verify that the current Symmetrix configuration is a viable configuration for making changes.
- ◆ Before creating new Symmetrix volumes, check for free physical disk space. New Symmetrix volumes are created first on physical disks that have no prior allocations, causing these disks to be committed to that emulation type.
- ◆ No configuration change is activated in the Symmetrix system until you commit the action.
- ◆ Some classes of change operations may or may not impact current I/O. When possible, before you commit any action, stop I/O activity on the Symmetrix volumes to be altered during a configuration change session.
- ◆ Ensure that all your critical data is preserved and safe when creating new or changing volume configurations. Do not store data on any volume that is not mirrored.
- ◆ After committing a mapping operation, you must update the volume mapping information within the host system environment. Attempting host activity with a volume after it has been removed or altered, but before you have updated the host's volume information, can cause host errors.

If I/O activity on an affected Symmetrix volume occurs before or during a commit action, the commit action might fail. At the very least, heavy I/O activity on unaffected volumes impacts how long it takes to commit changes.

Contact the EMC Customer Service for assistance in reverting to your previous configuration should there be unforeseen problems with the new configuration.

Understanding task persistence

Active EMC® Unisphere for VMAX configuration tasks persists across a server shutdown and subsequent restart. Inactive configuration tasks do not persist.

When the SMAS server is restarted, it restores all persisted tasks the task list, based on the user and the Symmetrix system.

The status of each restored task is determined by its status prior to the server shutdown, as detailed in the following table:

Table 2: Task status before and after server shutdown

Task status prior to server shutdown	Task status post server shutdown
Created	
Aborted	
Validated	
Running	
Successfully	

Task status prior to server shutdown	Task status post server shutdown
Run has error Prepare has error Submit has error Define has error Validate has error Done Failed	
Prepare in progress Abort in progress Submitted Prepared Defined Pending NA Acquiring lock Lock acquire failed Running During run During preview	Created
Commit in progress	Unknown ¹

1. Configuration tasks that were in a Commit in progress state prior to server shutdown will be restored in an Unknown state, since there is no way for the server to determine whether the task completed successfully. Therefore, to ensure that you are not committing a duplicate task, you should attempt to determine whether the original task completed successfully, prior to recommitting it.

Previewing/Running jobs

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.

Scheduling jobs

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job and click **View Details** to open its **Details** view.
4. Click **Schedule** to open the **Schedule** dialog box.
5. Select an **Execution Date** and **Execution Time**.
6. Click **OK**.

Running jobs

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job and click **Run**.

Rescheduling jobs

This procedure explains how to reschedule a job.

Before you begin:

- ◆ You cannot reschedule a single task in a job, only the whole job.
- ◆ If the volumes involved in a configuration change were originally reserved, rescheduling the tasks without reservation will not release the reservations. You need to manually release the reservations.

To reschedule a job:

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job, and click **View Details** to open its **Details** view.
4. Click **Schedule** to open the **Schedule** dialog box.
5. Select a new **Execution Date** and **Execution Time**.
6. Click **OK**.

Renaming jobs

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job, and click **View Details** to open its **Details** view.
4. Type a new name over the existing, and click **Apply**. Job names must be unique from other jobs on the Symmetrix system.

Reordering tasks within a job

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job, and click **View Details** to open its **Details** view.
4. In the task list table, select the task, and click **Move Up** or **Move Down**.

Grouping jobs

This procedure explains how to group two or more jobs into one job.

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select two or more jobs, and click **Group** to open the **Group Jobs** dialog box.
4. Specify a **Name for the New Job**, by doing one of the following:
 - ◆ Typing a name for the new job. Job names must be unique from other jobs on the Symmetrix system.
 - ◆ Selecting the name of one of the jobs to be grouped.
 - ◆ Use the default name, which is the next available short job ID selected by EMC® Unisphere for VMAX.
5. Click **OK**.

Un-grouping jobs

When un-grouping a job, EMC® Unisphere for VMAX will create an individual job for each of the tasks in the original job.

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job and click **View Details** to open its **Details** view.
4. In the task list, select **Remove** for each of the tasks jobs you want to un-group from the job.
5. Click **Apply**.

Stopping jobs

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select an Active job and click **Stop**.

If EMC® Unisphere for VMAX successfully stops the job, the state of the job change to **Stopped**.

Deleting jobs



1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select one or more jobs and click **Delete**.
4. Click **OK** in the confirmation message.
EMC® Unisphere for VMAX will remove the jobs from the job list view.

Viewing the job list

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.

The **Job List** view allows you to view and manage a job list for a Symmetrix system.

The following properties display:

- ◆ **Name** — User-defined name for the job or an ID assigned by EMC® Unisphere for VMAX, and an icon indicating if the job contains one () or more () tasks.
- ◆ **Status** — Status of the job. Possible values are:
 - **Scheduled** — Job is scheduled for execution.
 - **Created** — Job is created.
 - **Running** — Job is running. For jobs containing multiple tasks, this field will also display the job's progress. For example. Running (2 of 10).
 - **Completed** — Job finished running. This status displays for both succeeded and failed jobs.
 - **Pending** — Job is pending on the completion of another job.
 - **Stopped** — Job was running and a user stopped it.
 - **Missed** — Job was scheduled, but what was never run.
 - **Error** — Job failed.
- ◆ **User Name** — Host from which the job was created and the ID of user who created it.
- ◆ **Last Modified Time** — Date and time the job was moved to the job list.
- ◆ **Scheduled Time** — Date and time the job is scheduled to run.
- ◆ **Completed Time** — Date and time the job completed This field is blank for incomplete or failed jobs.

The following controls are available:

- ◆ **Run** — [Running jobs](#) on page 100
- ◆ **View Details** — [Viewing job details](#) on page 104
- ◆ **Group** — [Grouping jobs](#) on previous page
- ◆ **Stop** — [Stopping jobs](#) on previous page

- ◆ **Delete** — [Deleting jobs](#) on previous page

Viewing job details

1. Select the Symmetrix system.
2. Select **System** > **Job List** to open the **Job List** view.
3. Select a job and click **View Details** to open its **Details** view.

Use the job **Details** view to view and manage a job.

The following properties are displayed:

- ◆ **Name** — User-defined name for the job or an ID assigned by EMC® Unisphere for VMAX. To rename the job, type a new name over the existing and click **Apply**. Job names must be unique from other jobs on the Symmetrix system.
- ◆ **Status** — Status of the job. Possible values are:
 - **Scheduled** — Job is scheduled for execution.
 - **Created** — Job is created.
 - **Running** — Job is running. For jobs containing multiple tasks, this field will also display the job's progress. For example. Running (2 of 10).
 - **Completed** — Job finished running. This status displays for both succeeded and failed jobs.
 - **Pending** — Job is pending on the completion of another job.
 - **Stopped** — Job was running/appending and a user stopped it.
 - **Missed** — Job was scheduled, but was never run.
 - **Error** — Job failed.
- ◆ **Owner Name** — Host from which the job was created and the ID of user who created it.
- ◆ **Last Modified Time** — Date and time the job was moved to the job list.
- ◆ **Scheduled Time** — Date and time the job is scheduled to run.
- ◆ **Completed Time** — Date and time the job completed This field is blank for incomplete or failed jobs.
- ◆ **Task list** — Lists the configuration tasks in the job.
 - **Execution Order** — Order in which the task will be executed in relation to the other tasks in the job.
 - **Task** — Description of the task.
 - **Move Up/Move Down** — Allows you to reorder the tasks in the job list by moving a selected task up/down.

The following controls are available:

- ◆ **Run** — See [Running jobs](#) on page 100.
- ◆ **Schedule** — See [Scheduling jobs](#) on page 100.
- ◆ **Delete** — See [Deleting jobs](#) on page 102.
- ◆ **Stop** — See [Stopping jobs](#) on page 101.

- ◆ **Ungroup** — See [Un-grouping jobs](#) on page 101.
- ◆ **Apply** — Applies changes made in the job. For example, renaming the job.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Electronic licenses

Understanding licenses

EMC® Unisphere for VMAX supports Electronic Licensing (eLicensing) introduced with Enginuity 5875. eLicensing is an end-to-end license management solution to help you track and comply with software license entitlement. eLicensing leverages 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. This ensures that when upgrades are performed from a Symmetrix VMAX Family system running Enginuity versions lower than 5875 to a system running Enginuity 5875 or higher, the Symmetrix VMAX Family system is scanned for Enginuity features currently in use that require eLicenses. If Enginuity features are found in use, and there are no eLicenses registered and applied to support their use, they are internally reported as “IN USE,” which allows continued access to the Enginuity features while reporting that these features require proper licensing to ensure compliance. By only reporting this information, it prevents disruption to normal operations of your system and business. If your eLicensing report does display one or more Enginuity features as “IN USE,” it is your responsibility to work with your EMC Sales team to obtain proper eLicensing for those features.

With the introduction of eLicensing, Symmetrix licensing moved from a host-based model to a Symmetrix-based model, with the majority of licenses being stored internally on the Symmetrix system.

When installing licenses with eLicensing, you obtain license files from Powerlink, copy them to a Solutions Enabler or a EMC® Unisphere for VMAX host, and push them out to Symmetrix systems. Each license file fully defines all of 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, you must obtain a new license file from Powerlink and push it out to the Symmetrix system.

When managing your licenses, Solutions Enabler, EMC® Unisphere for VMAX, EMC z/OS Storage Manager (EzSM), MF SCF native command line, TPF, and IBM i platform console, allow you to view detailed usage reports so that you can better manage your capacity and compliance planning.

Licenses

Most Symmetrix licenses use the Symmetrix-based model (see [Symmetrix-based licenses](#) on page 111). However, there are still a number of Symmetrix licenses that remain host-based (see [Host-based licenses](#) on page 111). In addition, there are a number of host-based licenses being retired. The features that required the retired licenses still exist, they just no longer require licenses. For more on the licenses being retired with Enginuity 5875, refer to the *EMC Solutions Enabler Installation Guide*.

EMC® Unisphere for VMAX allows you to add and view Symmetrix-based licenses, and add, view, and remove host-based licenses.

Unisphere for VMAX licensing

Unisphere for VMAX uses Symmetrix-based eLicensing.

As a result, you can only manage a Symmetrix VMAX Family system from a Unisphere host, if the Symmetrix system contains a Unisphere for VMAX (SMC) eLicense. However, you can use Unisphere for VMAX to obtain and install the proper eLicense on the Symmetrix system.

Note Symmetrix VMAX family systems that do not contain a EMC® Unisphere for VMAX (SMC) eLicense will display a **Not licensed** link in the **Home** section of the EMC® Unisphere for VMAX interface. To install a EMC® Unisphere for VMAX (SMC) eLicense on such a Symmetrix system, click the link and follow the instructions in [Installing licenses](#) below.

When a Symmetrix VMAX Family system that has been managed by SMC is upgraded to Enginuity 5876 or higher, Unisphere will operate with the Symmetrix system even if the proper eLicense is not present. The Unisphere for VMAX (SMC) eLicense will be noted as *In Use* in Unisphere's eLicensing report. This designation means that the required eLicense is missing, but access to the system is still allowed to avoid service disruption. You should obtain and apply the proper eLicense to clear this designation.

Note With the incorporation of Symmetrix Performance Analyzer into Unisphere for VMAX (as the performance option), the former Symmetrix Performance Analyzer host-based eLicense is no longer required.

Note For more information on eLicensing, refer to the *EMC Solutions Enabler Installation Guide*.

Installing licenses

This procedure explains how to install Symmetrix and host-based licenses.

Before you begin:

- ◆ To perform this operation, you must be an Administrator or StorageAdmin.
- ◆ Before you can use EMC® Unisphere for VMAX to manage a Symmetrix system, you must first install a EMC® Unisphere for VMAX license on the system.
- ◆ To obtain Symmetrix-based licenses from Powerlink you will need the License Authorization Code (LAC) identification number from the LAC letter emailed to you.

To install licenses:

1. Select the Symmetrix system.
2. Select **System** › **Settings** › **Electronic Licenses** to open the **Electronic License** page.

3. Do the following, depending on the license type:
Symmetrix-based licenses:
 - a. Click **Symmetrix Entitlements** to open the **Symmetrix Entitlements** list view.
 - b. Click **Get File** to open the EMC Powerlink website, from which you can obtain new license files (by downloading or through email). Follow the instructions on the website. Be sure to download or copy the license file to a location on the SMAS server.
 - c. Click **Load File** to open the **Load License File** dialog box.
 - d. Click **Browse** and select the license file.
 - e. *Optional:* Click **Preview File** to preview the license file before loading it on the Symmetrix system.
 - f. Click **OK** to load the license file on the Symmetrix system.**Host-based licenses:**
 - a. Click **Solutions Enabler** to open the **Solutions Enabler License Key** list view.
 - b. Click **Load Key** to open the **Load License Key** dialog box.
 - c. Type the **License Key**.
 - d. Click **Apply** to load additional license keys. Click **OK** when done.

Removing host-based licenses

1. Select the Symmetrix system.
2. Select **System > Settings > Electronic Licenses > Solutions Enabler** to open the **Solutions Enabler** license view.
3. Select the license and click **Remove**.

Viewing Symmetrix entitlements

1. Select the Symmetrix system.
2. Select **System > Settings > Electronic Licenses** to open the **Electronic License** page.
3. Click **Symmetrix Entitlements** to open the **Symmetrix Entitlements** list view.

The **Symmetrix Entitlement** list view allows you view all licensed features (including those licensed with Symmetrix-based entitlements and host-based keys, which are still required) on Symmetrix systems running Engenuity 5875 or higher. In addition, this view also allows you obtain and install license files on Symmetrix systems running Engenuity 5875 or higher.

The following properties display:

- ◆ **Feature Name** — Name of the feature.
- ◆ **Activation Type** — Whether the feature's license is Permanent or an Evaluation copy. Evaluation licenses include an expiration date for reporting purposes only; the product title can still be used. Permanent licenses can be assigned to individual Symmetrix systems or to all the Symmetrix systems in the enterprise.

- ◆ **Activation Domain Type** — Whether the permanent license is assigned to an individual Symmetrix systems or to all the Symmetrix systems in the Enterprise. This column will appear blank for Evaluation type licenses.
- ◆ **Activation ID** — Activation ID assigned to the license file.
- ◆ **License** — Whether the license is host-based (**SE**) or Symmetrix-based (**EMCLM**).
- ◆ **Capacity Type** — Qualifies the capacity licensed. Possible values are:
 - **R-TB-Non-SATA** — Indicates that the capacity licensed applies to the raw capacity of all volumes on the system, excluding SATA.
 - **R-TB-SATA** — Indicates that the capacity licensed applies to the raw capacity of all SATA volumes on the system.
 - **REG-TB** — Indicates that the capacity licensed applies to the registered capacity of the Symmetrix system.
 - **R-TB External** — Indicates that the capacity licensed applies to the raw capacity of the virtualized LUNs in external storage.
- ◆ **Capacity** — Capacity licensed. The maximum quantity of data for which the functionality of the software is licensed to use, in Terabytes.
- ◆ **Status** — Whether the license is Enabled or Disabled.
- ◆ **Expiration Date** — Date an evaluation license expires. For permanent licenses, this field appears blank.
- ◆ **Install Date** — Date the license file was installed on the Symmetrix system.

The following controls are available:

- ◆ **View File** — Opens the actual license file installed on the Symmetrix system.
- ◆ **Get File** — Opens the EMC Powerlink website, from which you can obtain a license file (by downloading or through email). You must download or copy a license file to the SMAS server before you can apply it to a Symmetrix system.
- ◆ **Load File** — See [Installing licenses](#) on page 107.

Viewing host-based licenses

1. Select the Symmetrix system.
2. Select **System > Settings > Electronic Licenses > Solutions Enabler** to open the **Solutions Enabler** license view.

The **Solutions Enabler** license view allows you to view and manage host-based licenses.

The following properties display:

- ◆ **License Key** — License key.
- ◆ **Licensed Features** — Name of the licensed feature.

The following controls are available:

- ◆ **Load Key** — See [Installing licenses](#) on page 107.
- ◆ **Remove Key** — See [Removing host-based licenses](#) on page 108.

Viewing license usage

1. Select the Symmetrix system.
2. Select **System > Settings > Electronic Licenses > Show License Usage** to open the **License Usage** view.

The **License Usage** view shows the results of a query to the Symmetrix system's feature registration database (Enginuity 5875 or higher).

The following properties display:

- ◆ **Feature Name** — Name of the feature.
- ◆ **Activation Type** — How the product title was activated. Possible values are:
 - **Entitlement** — Indicates that product title was activated through an entitlement.
 - **Manual** — Indicates that the product title was manually activated by EMC.
 - **In Use** — Indicates that the product title is activated because it was in use prior to upgrading from Enginuity 5874 to Enginuity 5875.
Product titles activated manually (MAN) or because they were in use (USE) are not considered properly entitled, in which case you should contact EMC for proper entitlement.
- ◆ **Capacity Type** — qualifies the licensed capacity . Possible values are:
 - **R-TB-Non-SATA** — Indicates that the capacity licensed applies to the raw capacity of all volumes on the system, excluding SATA.
 - **R-TB-SATA** — Indicates that the capacity licensed applies to the raw capacity of all SATA volumes on the system.
 - **REG-TB** — Indicates that the capacity licensed applies to the registered capacity of the Symmetrix system.
 - **R-TB External** — Indicates that the capacity licensed applies to the raw capacity of the virtualized LUNs in external storage.

- ◆ **Licensed Capacity** — Maximum quantity of data for which the functionality of the software is licensed to use, in Terabytes.
- ◆ **Used Capacity** — Amount of licensed capacity currently in use.

Host-based licenses

The following host-based licenses remain unchanged, regardless of Engenuity version.

License	Description
FAST for DMX	Cross Disk Group
TimeFinder	TimeFinder/Mirror

Symmetrix-based licenses

With the release of Engenuity 5876, EMC® Unisphere for VMAX supports Symmetrix-based license bundles. A license bundle is a single license that enables multiple features. For example, the Symmetrix Remote Replication Suite license bundle enables the SRDF, SRDF/A and SRDF/S product titles. For Symmetrix systems running versions of Engenuity 5875, EMC® Unisphere for VMAX continues to support individual Symmetrix-based licenses.

Table 3: Symmetrix Licenses

License/Description		Allows you to
Engenuity 5876 or higher	Engenuity 5875 ¹	
SYMM_VMAX_Engenuity License for whole system Includes: Dynamic Cache Partitioning Symmetrix Priority Controls Symmetrix Optimizer	SYMM_Model_Engenuity License for whole system	Virtualize an e Disk for encapsulation
		Use VLUN to migrate from an encapsulated volume (use it as a source volume)
		Use an encapsulated volume as a clone source
	SYMM_Model_DCP ² Dynamic Cache Partitioning	Enable cache partitions for a Symmetrix system
		Create cache partitions
		Set cache partitions to Analyze mode
	SYMM_Model_SPC ² Symmetrix Priority Controls	Enable priority of service for a Symmetrix system
		Set host I/O priority
		Set copy QoS priority

License/Description		Allows you to
Enginuity 5876 or higher	Enginuity 5875 ¹	
	SYMM_Model_OPTIMIZER ² Symmetrix Optimizer	Enable Optimizer functionality, including: <ul style="list-style-type: none"> ◆ Manual mode ◆ Rollback mode ◆ Manual Migration mode
		Schedule manual swaps
		Set the following Optimizer-specific parameters: <ul style="list-style-type: none"> ◆ Device Swap Priority ◆ Any of the Optimizer Advanced parameters
		Set the following Optimizer/FAST parameters: <ul style="list-style-type: none"> ◆ User Approval Mode ◆ Maximum Devices to Move ◆ Maximum Simultaneous Devices ◆ Workload Period ◆ Minimum Performance Period
		Validate or create VLUN migrations
		Create time window

License/Description		Allows you to
Engenuity 5876 or higher	Engenuity 5875 ¹	
SYMM_VMAX_SRDF_REPLICATION Symmetrix Remote Replication Suite Includes: SRDF SRDF/Asynchronous mode SRDF/Synchronous mode	SYMM_Model_SRDF ^{2, 3} SRDF	Create new RDF groups

License/Description		Allows you to
Engenuity 5876 or higher	Engenuity 5875 ¹	
		Create dynamic RDF pairs in Adaptive Copy mode
		Create RDF volumes
		Convert non_RDF volumes to RDF
		Add RDF mirrors to volumes in Adaptive Copy Mode
		Set the dynamic-RDF capable attribute on volumes
		Create SAVE volumes
	SYMM_Model_SRDF_A ² SRDF/Asynchronous mode	Create dynamic RDF pairs in Asynchronous mode
		Set RDF pairs into Asynchronous mode
		Add RDF mirrors to volumes in Asynchronous mode
		Create RDFA_DSE pools
		Set any of the following SRDF/A attributes on an RDF group: <ul style="list-style-type: none"> ◆ Minimum Cycle Time ◆ Transmit Idle ◆ DSE attributes, including: <ul style="list-style-type: none"> • Associating an RDFA-DSE pool with an RDF group • DSE Thresholds • DSE Autostart ◆ Write Pacing attributes, including: <ul style="list-style-type: none"> • Write Pacing Thresholds • Write Pacing Autostart • Device Write Pacing exemption • TF Write Pacing Autostart

License/Description		Allows you to
Engenuity 5876 or higher	Engenuity 5875 ¹	
	SYMM_Model_SRDF_S ² SRDF/Synchronous mode	Create dynamic RDF pairs in Synchronous mode Set SRDF pairs in Synchronous mode Add an RDF mirror to a volume in Synchronous mode
SYMM_VMAX_SRDF_STAR SRDF/Star	SYMM_Model_SRDF_STAR ² SRDF/Star	Perform a setup to utilize the environment
SYMMETRIX_VMAX_TIMEFINDER Symmetrix TimeFinder Suite Includes: TimeFinder/Clone TimeFinder/Snap	SYMM_Model_TF_CLONE ² TimeFinder/Clone	Create new native clone sessions
		Create new TimeFinder/Clone emulations
	SYMM_Model_TF_SNAP ² TimeFinder/Snap	Create new sessions
		Duplicate existing sessions
		Create snap pools
		Create SAVE volumes

License/Description		Allows you to
Enginuity 5876 or higher	Enginuity 5875 ¹	
SYMM_VMAX_FAST_TIERING Symmetrix Tiering Suite Includes: FAST for disk groups FAST for virtual pools	SYMM_Model_FAST FAST for disk groups	Create time windows
		Add disk group tiers to FAST Policies
		Enable FAST
		Set the following Optimizer/FAST parameters: <ul style="list-style-type: none"> ◆ Swap Non-Visible Devices ◆ Allow Only Swap⁴ ◆ User Approval Mode ◆ Maximum Devices to Move ◆ Maximum Simultaneous Devices ◆ Workload Period ◆ Minimum Performance Period
	SYMM_Model_FAST_VP FAST for virtual pools	Create time windows
		Add virtual pool (VP) tiers to FAST Policies
		Enable FAST
		Set the following FAST VP-specific parameters: <ul style="list-style-type: none"> ◆ Thin Data Move Mode ◆ Thin Relocation Rate ◆ Pool Reservation Capacity
	Set the following Optimizer/FAST parameters: <ul style="list-style-type: none"> ◆ Workload Period ◆ Minimum Performance Period 	

License/Description		Allows you to
Enginuity 5876 or higher	Enginuity 5875 ¹	
SYMM_VMAX_OR_DM RCOPY	SYMM_Model_OR_DM ² RCOPY	Create host push sessions
		Create cold pull sessions
		Create cold push sessions
SYMM_VMAX_SMC	SYMM_Model_SMC Symmetrix Management Console	Manage Symmetrix systems running Enginuity 5875 Q2 2011 SR or higher
SYMM_VMAX_FTS Federated Tiered Storage	SYMM_Model_FTS ⁵ Federated Tiered Storage	Virtualize an eDisk for external provisioning

¹In the license name, *Model* indicates the Symmetrix model on which the license is installed. Possible values are: VMAX and VMAXe.

²This feature requires a host-based license when the Symmetrix system being managed is running an Enginuity version lower than 5875.

³Requires that the Symmetrix system also be licensed for SRDSF/A (SYMM_Model_SRDF_A) and/or SRDF/S (SYMM_Model_SRDF_S).

⁴Requires that the Symmetrix system also be licensed for Optimizer (SYMM_Model_OPTIMZER)

⁵This product title is not supported on Symmetrix VMAX10K systems.

Symmetrix access controls

Understanding access controls

Administrators , **StorageAdmins**, and **SecurityAdmins** can set access controls on specific volumes within a Symmetrix 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 connected to that Symmetrix system will not see those volumes. This eliminates the possibility of one host inadvertently performing operations on volumes that belong to someone else.

To set up access controls:

1. Create an access control group. See [Creating access groups](#) on the facing page..
2. Add to the group, one or more hosts (access name) and unique IDs. Each host has a unique ID; a group can contain one or more hosts. See [Adding access ID to access groups](#) on the facing page..
3. Create one or more pools of volumes. Specific volumes can belong to only one pool. See [Creating access pools](#) on page 121.
4. Create one or more access control entries. An access control entry associates a pool with a group, and grants the access control types. See [Creating access control entries](#) on page 123..

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

Opening access controls

Access control dialogs and views are restricted and require you to enter a pin provided by EMC Customer Service.

Before you begin:

To perform this operation, you must be an Administrator or SecurityAdmin.

To open access controls:

1. Select a Symmetrix system.
2. Select **System** > **Settings** > **Symmetrix Access Controls** to open the **Access Control Pin Entry** dialog box.
3. Do one of the following:
 - ◆ Select **Read Only** to gain read only access to the **Symmetrix Access Controls** page.
 - ◆ Type the **Access Control Pin**.
4. Click **OK**.

Creating access groups

Typically various sets of users tend to use the same applications that utilize common EMC® Unisphere for VMAX features from a given host. They typically require the same volume resources and permissions of access to these shared volumes. For this reason, hosts are registered in groups identified with a group name, which serves as a root for all ACEs in the group.

Access groups contain groups of access IDs and their ID names. Any ID and name must belong to just one group and are entered into the database together. For ease of management, it is highly recommended that you choose an access ID name that best associates with the particular host in use. For example, SunHost1 is more appropriate than a name such as JRSMITH.

Once the group is created, the group name can be used to create access control entries (ACEs).

To create an access control group:

1. Select a Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.
5. Click **Create** to open the **Create Access ID Group** dialog box.
6. Type a **Group Name**. Access group names must be unique from other access pools on the system and cannot exceed 8 characters (Engenuity 5771 and lower) or 31 characters (Engenuity 5772 or higher). Only alphanumeric characters, underscores (_), and dashes (-) are allowed. Access group names are case-sensitive.
7. Type the **Host ID**. This value is the host computer name.
8. Type the host's **Unique Access ID**. To find this value, run the SYMCLI command `symacl - unique` on the host computer, (host ID example: 2C5E05B6-53408AC9-9C3F747C).
9. Click **Add Host**. This value is the name of the access ID (eight characters).
10. Click **OK**.

Adding access ID to access groups

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.
5. Select the group and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Access IDs** to open the **Access ID** list view.
7. Click **Add** to open the **Add Access ID** dialog box.
8. To add host access IDs to the group:
 - a. Type the **Host ID**. This value is the host computer name.
 - b. Type the host's **Unique Access ID**. To find this value, run the SYMCLI command `symacl - unique` on the host computer, (host ID example: 2C5E05B6-53408AC9-9C3F747C).
 - c. Click **Add Host**.

To add user access IDs to the group:

- a. Type the **User ID name**. This value is the ID assigned by the access control administrator, must be between four and twelve characters long.
 - b. Type the user's access control pin.
 - c. Click **Add User**.
9. Click **OK**.

Removing access IDs from access groups

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.
5. Select the group and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Access IDs** to open the **Access ID** list view.
7. Select one or more IDs and click **Remove**.
8. Click **OK** in the confirmation message.

Deleting access groups

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.
5. Select one or more groups and click **Delete**.
6. Click **OK** in the confirmation message.

Viewing access groups

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.

Use the **Access Groups** list view to display and manage access groups on the Symmetrix system.

The following properties display:

- ◆ **Access Group ID** — Name of the access ID.
- ◆ **Number of Access IDs** — Number of members (host access IDs) in the group.
- ◆ **Number of Access Control Entries** — Number of access pools associated with the group.

The following controls are available:

- ◆ **Create** — See [Creating access groups](#) on previous page.

- ◆ **View Details** — See [Viewing access group details](#) on page 125.
- ◆ **Delete** — See [Deleting access groups](#) on previous page.

Creating access pools

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.
5. Click **Create** to open the **Create Access Controlled Volume Pool** dialog box.
6. Type a **Pool Name**. Access pool names must be unique from other access pools on the system and cannot exceed 8 characters (Engenuity 5771 and lower) or 31 characters (Engenuity 5772 or higher). Only alphanumeric characters, underscores (_), and dashes (-) are allowed. Access pool names are case-sensitive.
7. Select one or more **Available Volumes** and click **Add Volume** to move them to the **Selected Volumes** list.

Note When adding a meta volumes to an access pool, you must select the meta head. Individual meta members do not appear in the volume lists.
8. Click **OK**.

Modifying access pools

The following explains how to add/remove volumes from an access pool:

1. Select a Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.
5. Select the access pool and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Volumes** to the access pool **Volumes** list view.
7. To add volumes to the pool:
 - a. Click **Add** to open the **Add Volume to Access Pool** dialog box.
 - b. Select one or more **Available Volumes** and click **Add Volume** to move them to the **Selected Volumes** list.
 - c. Click **OK**.

To remove volumes from the pool:

 - a. Select one or more volumes and click **Remove**.
 - a. Click **OK** in the confirmation message.

Deleting access pools

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.

3. Type the **Access Control Pin** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.
5. Select one or more pools and click **Delete**.
6. Click **OK** in the confirmation dialog box.

Viewing access pools

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.

Use the **Access Pools** list view to display and manage Access Pools on the Symmetrix system.

The following properties display

- ◆ **Access Controlled Volume Pool** — Name of the pool.
- ◆ **Number of Volumes** — Number of volumes in the pool.
- ◆ **Number of Access Controlled Entries** — Number of access groups associated with the pool.

The following controls are available:

- ◆ **Create** — See [Creating access pools](#) on previous page.
- ◆ **View Details** — See [Viewing access pool details](#) on page 126.
- ◆ **Delete** — See [Deleting access pools](#) on previous page.

Modifying access types

The following explains how to add/remove access types defined in an access control entry.

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **ACE** to open the **ACEs** list view.
5. Select the entry and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Access Types** to open the **Access Types** list view.
7. To add access types:
 - a. Click **Add** to open the **Assign Access Types** dialog box.
 - b. Select one or more access types and click **Add**. For more on the available access types, see [Access types](#) on page 128..
 - c. Click **OK**.

To remove access types:

- a. Select one or more and click **Remove**.
- b. Click **OK** in the confirmation message.

Creating access control entries

1. Select the Symmetrix system.
2. Select **System** › **Settings** › **Symmetrix Access Controls**.
3. Type the **Access Control Pin**.
4. Click **ACE** to open the **ACEs** list view.
5. Click **Create** to open the **Create Access Control Entries** dialog box.
6. Select/Create the access group to include in the ACE.
7. Select/Create the access pool to include in the ACE.
8. Select one or more **Available Access Types** and click **Add Access Type** to move them to the **Selected Access Types**.
9. Click **OK**.

Deleting access control entries

1. Select the Symmetrix system.
2. Select **System** › **Settings** › **Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **ACE** to open the **ACEs** list view.
5. Select one or more entries and click **Remove**.
6. Click **OK** in the confirmation message.

Viewing access control entries

1. Select the Symmetrix system.
2. Select **System** › **Settings** › **Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **ACE** to open the **ACEs** list view.

Use the **ACEs** list view to display and manage access control entries on the Symmetrix system.

The following properties display:

- ◆ **Access ID Group** — Associated access group.
- ◆ **Access-controlled Volume Pool** — Associated access pool.
- ◆ **Access Type** — .Permissions assigned to the group/pool. For possible values, see See [Access types](#) on page 128..

The following controls are available:

- ◆ **Create** — See [Creating access control entries](#) above.
- ◆ **View Details** — See [Viewing access control entry details](#) on next page.
- ◆ **Delete** — See [Deleting access control entries](#) above.

Viewing access control entry details

1. Select the Symmetrix system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **ACE** to open the **ACEs** list view.
5. Select an entry and click **View Details** to open its **Details** view.

Use the **ACE Details** view to display and an access control entry. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Access ID Group** — Associated access group.
- ◆ **Access-controlled Volume Pool** — Associated access pool.
- ◆ **Access Type** — .Permissions assigned to the group/pool. For possible values, see [See Access types](#) on page 128..

The following controls are available:

- ◆ **Create** — See [Creating access control entries](#) on previous page.
- ◆ **Delete** — See [Deleting access control entries](#) on previous page.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the access control entry. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Access IDs - 2** will open a view listing the two access IDs associated with the entry.

Viewing access groups

1. Select the Symmetrix system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.

Use the **Access Groups** list view to display and manage access groups on the Symmetrix system.

The following properties display:

- ◆ **Access Group ID** — Name of the access ID.
- ◆ **Number of Access IDs** — Number of members (host access IDs) in the group.
- ◆ **Number of Access Control Entries** — Number of access pools associated with the group.

The following controls are available:

- ◆ **Create** — See [Creating access groups](#) on page 119.
- ◆ **View Details** — See [Viewing access group details](#) below.
- ◆ **Delete** — See [Deleting access groups](#) on page 120.

Viewing access group details

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.
5. Select the group and click **View Details** to open its **Details** view.

Use the access group **Details** view to display and manage an access group. This view contains two panels, [Properties](#), and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Access Group ID** — Name of the access ID.
- ◆ **Number of Access IDs** — Number of members (host access IDs) in the group.
- ◆ **Number of Access Control Entries** — Number of access control pools associated with the group.

The following controls are available:

- ◆ **Create** — See [Creating access groups](#) on page 119.
- ◆ **Delete** — See [Deleting access groups](#) on page 120.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the access group. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Access IDs - 2** will open a view listing the two access IDs in the group.

Viewing access IDs

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Groups** to open the **Access Groups** list view.
5. Select the group and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Access IDs** to open the **Access ID** list view.

Use the **Access ID** list view to display and manage a access IDs.

The following property displays:

- ◆ **Access ID** — Access IDs assigned to the group.

The following controls are available:

- ◆ **Add** — See [Adding access ID to access groups](#) on page 119.
- ◆ **Remove** — See [Removing access IDs from access groups](#) on page 120.

Viewing access pools

1. Select the Symmetrix system.
2. Select **System** > **Settings** > **Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.

Use the **Access Pools** list view to display and manage Access Pools on the Symmetrix system.

The following properties display

- ◆ **Access Controlled Volume Pool** — Name of the pool.
- ◆ **Number of Volumes** — Number of volumes in the pool.
- ◆ **Number of Access Controlled Entries** — Number of access groups associated with the pool.

The following controls are available:

- ◆ **Create** — See [Creating access pools](#) on page 121.
- ◆ **View Details** — See [Viewing access pool details](#) below.
- ◆ **Delete** — See [Deleting access pools](#) on page 121.

Viewing access pool details

1. Select the Symmetrix system.
2. Select **System** > **Settings** > **Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.
5. Select the access pool and click **View Details** to open its **Details** view.

Use the access pool **Details** view to display and manage an access pool. This view contains two panels, [Properties](#), and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Access Controlled Volume Pool** — Name of the pool.
- ◆ **Number of Volumes** — Number of volumes in the pool.
- ◆ **Number of Access Controlled Entries** — Number of access groups associated with the pool.

The following controls are available:

- ◆ **Create** — See [Creating access pools](#) on page 121.
- ◆ **Delete** — See [Deleting access pools](#) on page 121.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the access pool. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Volumes - 2** will open a view listing the two volumes contained in the pool.

Viewing access pool volumes

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access Pools** to open the **Access Pools** list view.
5. Select the access pool and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Volumes** to the access pool **Volumes** list view.

Use the access pool **Volumes** list view to display and manage the volumes in an access pool.

The following properties display

- ◆ **Volume ID** — Volume identifier.

The following controls are available:

- ◆ **Add** — See [Modifying access pools](#) on page 121.
- ◆ **Remove** — See [Modifying access pools](#) on page 121.

Viewing access types

1. Select the Symmetrix system.
2. Select **System > Settings > Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **ACE** to open the **ACEs** list view.
5. Select an entry and click **View Details** to open its **Details** view.
6. In the **Related Objects** panel, click **Access Types** to open the **Access Types** list view.

Use the **Access Types** list view to display and manage the access types associated with an access control entry.

The following properties display

- ◆ **Access Type** — Permissions assigned to the group/pool. For possible values, see See [Access types](#) on next page...

The following controls are available:

- ◆ **Add** — [Adding access ID to access groups](#) on page 119
- ◆ **Remove** — [Removing access IDs from access groups](#) on page 120

Access types

This table lists/describes the possible access types. Access define the permissions assigned to access groups and pools.

Access type	Description
ADMIN	Grants administrator privilege to grant/deny access control entries to hosts and users.
ADMINRD	Grants read access only to all access control information.
ALL	All possible access types granted except ADMIN and ADMINRD. Must be directed to ALL volumes.
BASE	Allows the discovery of devices and to obtain states and statistics from the Symmetrix system (directors and volumes).
BASECTRL	Allows base control operations on volumes and device groups.
BCV	Allows TimeFinder (BCV) and clone control and status operations.
CACHCTRL	Allows cache control operations concerning LRU partition management.
CFGDEV	Allows powerful configuration control operations that manage various types of configuration changes on volumes in the Symmetrix.
CFGSYM	Allows access to set Symmetrix attributes, set port flags, and swap RA groups. Must be directed to ALL volumes.
CHECKSUM	Allows Symmetrix volume Double Checksum operations.
CREATEDV	Allows the creation and deletion of Symmetrix volumes.
DIRCTRL	Allows you to take directors and their ports offline and online. Must be directed to ALL volumes.
ECC	Allows the ECC Symmetrix agent to run on the requested host.
OPTMZR	Allows user-configurable attributes that may affect the Optimizer behavior.
POWRPATH	Access to PowerPath-directed devices in an RDF group. Must be directed to ALL volumes.
QOS	Allows the execution of Quality of Service (QOS) performance control operations to manage copy priorities. Excludes LRU cache control functionality.
RCOPY	Manages Open Replicator sessions.
RDF	Allows SRDF control and set operations.
SDDF	Allows the DeltaMark (Change Tracker) functionality that monitors track changes.
SDR	Allows mapping/unmapping of devices to directors/ports for the Symmetrix Disk Reallocation (SDR) feature.

Access type	Description
SNAP	Allows the creation and management of virtual copy sessions between a source volume and multiple virtual (VDEV) target volumes.
VLOGIX	Enables access to volume Masking or Volume Logix volumes.

Dynamic Cache Partitioning

Enabling/Disabling dynamic cache partitioning

1. Select a Symmetrix system.
2. Select **System** › **Settings** › **Symmetrix Attributes** to open the **Symmetrix Attributes** page.
3. In the **DCP** panel, set the **Cache Partition Status** to **Enable**.
4. Click **Apply**.

Creating dynamic cache partitions

This procedure explains how to create dynamic cache partitions.

Before you begin:

- ◆ There must be an available partition.
- ◆ There must be enough cache left in the default partition that it does not fall below the minimum required cache.
- ◆ The number of cache partitions allowed on a Symmetrix system is defined in the Symmetrix system's properties file. The maximum number allowed is 16.
- ◆ The sum of target % for all defined partitions must be 100%.


To create a dynamic cache partition:

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel click **Dynamic Cache Partitions** to open the **Dynamic Cache Partitions** list view.
4. Click **Create** to open the **Create Dynamic Cache Partition** dialog box.
5. Type a **Name** for the dynamic cache partition. Dynamic cache partition names must be unique from other cache partition names on the Symmetrix system and cannot exceed 32 characters. Only alphanumeric characters and underscores are allowed. Note that underscores can only be used with the string; not on the ends of the string.
6. Type the minimum target percentage (**Min Target %**), for the partition. This values must be less than the Target %.
7. Type the target cache percentage (**Target %**) for the partition. This value must be less than the Max Target %.
8. Type the maximum cache percentage (**Max Target %**) for the partition.
9. Type the **Donation Time** in seconds. This value is the length of time before idle cache will be made available to other partitions. The default value is 300 seconds.

10. Type the write pending limit percentage (**WP Limit %**) for the cache partition. Possible values 40-80, with 80 being the default.
11. Click **OK**.

Assigning volumes to dynamic cache partitions

This procedure explains how to assign volumes to dynamic cache partitions at the volume level. You can also perform this operation from the storage group level, the device group level, or the RDF group level.

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, select the type of volume.
4. Click **View** to open its list view.
5. Select one or more volumes, click more , and select **Assign Dynamic Cache Partition** to open the **Assign Dynamic Cache Partition** dialog box.
6. *Optional:* Click **Show Selected Volumes** to view details on the selected volumes.
7. Select a **Dynamic Cache Partition** and click **OK**.
8. Click **OK** in the confirmation message.

Deleting dynamic cache partitions

1. Select the Symmetrix system.
2. Select **System > Dashboard**.
3. In the **Hardware** panel, click **Dynamic Cache Partitions** to open the **Dynamic Cache Partitions** list view.
4. Select one or more partitions and click **Delete**.
5. Click **OK** in the confirmation message.

Running in analyze mode

Cache partitioning Analyze mode is a tool for helping you determine the amount of cache your applications are consuming, prior to enabling the cache partitioning feature. Once you have determined the amount of cache your applications are consuming, you can then modify the existing partitions, or add/delete partitions to achieve the required performance.

Enabling Analyze mode will automatically set the following cache partition settings:

- ◆ Max % = 100
- ◆ Min % = 0
- ◆ Donation Time = 0 (seconds)

These settings will allow cache to behave as if there are no partitions.

To use this tool:

1. Enable dynamic cache partitioning in Analyze mode.
2. Create your dynamic cache partitions. See [Creating dynamic cache partitions](#) on page 130.
3. Assign volumes to the cache partitions.
4. Monitor cache usage using the QoS Monitor.
5. Once you have gathered enough usage data, change the cache partitioning status from Analyze mode to Enable.
6. Make changes to the cache partitions based on the usage data.

Viewing dynamic cache partitions

1. Select the Symmetrix system.
2. Select **System** > **Dashboard**.
3. In the **Hardware** panel, click **Dynamic Cache Partitions** to open the **Dynamic Cache Partitions** list view.

Use the **Dynamic Cache Partitions** list view to display and manage dynamic cache partitions.

The following properties display:

- ◆ **Name** — Name of the partition.
- ◆ **Min %** — Minimum target percentage.
- ◆ **Tgt %** — Target cache percentage.
- ◆ **Max %** — Maximum cache percentage.
- ◆ **Donation Time** — Length of time before idle cache will be made available to other partitions.
- ◆ **WP Limit** — Write pending limit percentage
- ◆ **Slots Used** — Number of cache slots used by the partition.
- ◆ **% Used** — Percentage of cached used by the partition.

The following controls are available:

- ◆ **Create** — See [Creating dynamic cache partitions](#) on page 130.
- ◆ **View Details** — See [Viewing dynamic cache partition details](#) on the facing page.
- ◆ **Delete** — See [Deleting dynamic cache partitions](#) on previous page.

Viewing dynamic cache partition details

1. Select the Symmetrix system.
2. Select **System** › **Dashboard**.
3. In the **Hardware** panel, click **Dynamic Cache Partitions** to open the **Dynamic Cache Partitions** list view.
4. Select a partition and click **View Details** to open its **Details** view.

The cache partition **Details** view allows you to view and manage a cache partition. This view contains [Properties](#), [Related Objects](#), and [Performance Views](#) panels.

Properties panel

The following properties display:

- ◆ **Name** — Name of the partition. To rename the partition, type a new name over the existing and click **Apply**. Dynamic cache partition names must be unique from other cache partition names on the Symmetrix system and cannot exceed 32 characters. Only alphanumeric characters and underscores are allowed. Note that underscores can only be used with the string; not on the ends of the string. You cannot modify the name of the DEFAULT_PARTITION.
- ◆ **Min Target %** — Minimum target percentage. To change this value, type a new value over it and click **Apply**. This value must be less than the Target %. You cannot modify this value for the DEFAULT_PARTITION.
- ◆ **Target %** — Target cache percentage. To change this value, type a new value over it and click **Apply**. This value must be less than the Max Target %. You cannot modify this value for the DEFAULT_PARTITION.
- ◆ **Max Target %** — Maximum cache percentage. To change this value, type a new value over it and click **Apply**. You cannot modify this value for the DEFAULT_PARTITION.
- ◆ **Donation Time** — Length of time before idle cache will be made available to other partitions. To change this value, type a new value over it and click **Apply**.
- ◆ **Write Pending Limit** — Write pending limit percentage. Possible values 40-80, with 80 being the default.
- ◆ **Write Pending Slot Count** — Write pending slot count.
- ◆ **Cache Slots Used** — Number of cache slots used by the partition.
- ◆ **Cache Percentage Used** — Percentage of cache used by the partition.

The following controls are available:

- ◆ **Create** — See [Creating dynamic cache partitions](#) on page 130.
- ◆ **Delete** — See [Deleting dynamic cache partitions](#) on page 131.
- ◆ **Apply** — Applies changes made in the **Properties** panel.
- ◆ **Cancel** — Cancels changes made in the **Properties** panel.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the cache partition. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Volumes - 2** will open a view listing the two volumes contained in the cache partition.

Performance Views panel

The **Performance** panel links you to the performance analyze views for the group.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

CHAPTER 4

Storage Management


This chapter contains the following:

- Assigning Symmetrix priority 136
- Setting replication QoS 137
- Symmetrix VMAX provisioning 139
- Storage groups 148
- Fully Automated Storage Tiering 169
- Volume configuration 194
- Disk groups 221
- Virtual Provisioning 226
- Enhanced Virtual LUN migration 248
- Storage templates 252
- External storage 258
- Reservations 263
- Optimizer 266


Assigning Symmetrix priority

You can assign Symmetrix priority to individual volumes or groups of volumes (DGs or SGs).

To assign Symmetrix priority to individual volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, select the type of volume.
4. Click **View** to open its list view.
5. Select one or more volumes, click more , and select **Assign Symmetrix Priority** to open the **Assign Symmetrix Priority** dialog box.
6. *Optional:* Click **Show Selected Volumes** to view details on the selected volumes.
7. Select a **Symmetrix Priority** from 1 (the fastest) to 16 (the slowest) and click **OK**.
8. Click **OK** in the confirmation message.



To assign Symmetrix priority to groups of volumes:

1. Select the Symmetrix system.
2. Select the group (**Storage > Storage Groups**, or **Data Protection > Device Groups**).
3. Click more  and select **Assign Symmetrix Priority** to open the **Assign Symmetrix Priority** dialog box.
4. *Optional:* Click **Show Selected Volumes** to view details on the selected volumes.
5. Select a **Symmetrix Priority** from 1 (the fastest) to 16 (the slowest) and click **OK**.
6. Click **OK** in the confirmation message.

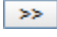
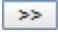
Setting replication QoS

The QoS (Quality of Service) feature, allows you to adjust the data transfer (copy) pace on individual volumes or groups of volumes (DGs or SGs) for certain operations. By increasing the response time for specific copy operations you can increase the overall performance of other Symmetrix volumes. You can set the copy pace for a [storage group/volumes](#) or a [device group/volumes](#).

To set the copy pace for storage groups:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Perform one of the following actions:
 - ◆ *All volumes in the storage group*: Select the storage group, click more  , and select **Replication QoS**.
 - ◆ *Some volumes in the storage group*: Select the storage group, click **View Details**, and select **Volumes - nn** from the **Related Objects** panel to open the **Volumes** list view. Select the volumes, click more  , and select **Replication QoS**.
4. Select the **Operation Type** from the possible values:
 - ◆ **BCV** — Sets the copy pace priority during BCV operations.
 - ◆ **RDF** — Sets the copy pace priority during RDF operations.
 - ◆ **Mirror Copy** — Sets the copy pace priority during mirror operations.
 - ◆ **Clone** — Sets the copy pace priority during Clone operations.
 - ◆ **VLUN** — Sets the copy pace priority during virtual LUN migrations.
5. Select the **Copy Pace** from the possible values:
 - ◆ **0-16** — Sets the copy pace, with 0 being the fastest (and the default) and 16 being the slowest.
 - ◆ **STOP** — Stops the copy. Not supported when the Operation Type is BCV, or the Symmetrix system is running an Enginuity version earlier than 5875.
 - ◆ **URGENT** — Sets the copy pace to urgent, which may be faster than the default (0). Not supported when the **Operation Type** is BCV, or the Symmetrix system is running an Enginuity version earlier than 5875.

To set the copy pace for device groups:

1. Select the Symmetrix system › **Data Protection** › **Device Groups** to open the **Device Groups** view.
2. Perform one of the following actions:
 - ◆ *All volumes in the device group:* Select the device group, click  **More**, and select **Replication QoS**.
 - ◆ *Some volumes in the device group:* Select the device group, click **View Details**, and select **Volumes - nn** from the Related Objects panel to open the **Volumes** view. Select the volume(s), click  **More**, and select **Replication QoS**.
3. Select the **Operation Type** from the possible values:
 - ◆ **BCV** — Sets the copy pace priority during BCV operations.
 - ◆ **RDF** — Sets the copy pace priority during RDF operations.
 - ◆ **Mirror Copy** — Sets the copy pace priority during mirror operations.
 - ◆ **Clone** — Sets the copy pace priority during Clone operations.
 - ◆ **VLUN** — Sets the copy pace priority during virtual LUN migrations.
4. Select the **Copy Pace** from the possible values:
 - ◆ **0-16** — Sets the copy pace, with 0 being the fastest (and the default) and 16 being the slowest.
 - ◆ **STOP** — Stops the copy. Not supported when the Operation Type is BCV, or the Symmetrix system is running an Enginuity version earlier than 5875.
 - ◆ **URGENT** — Sets the copy pace to urgent, which may be faster than the default (0). Not supported when the **Operation Type** is BCV, or the Symmetrix system is running an Enginuity version earlier than 5875.
5. *If performing this operation on a group:* Select one of the **Apply Change to** values: **All Devices**, **BCV Devices Only**, or **non-BCV Devices Only**.

Symmetrix VMAX provisioning

Provisioning Symmetrix VMAX storage overview

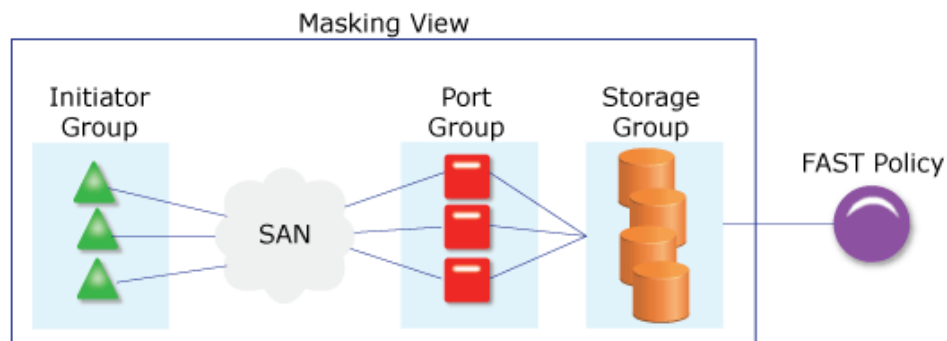
Provisioning storage refers to the process by which you make storage available to hosts.

EMC® Unisphere for VMAX provides the following methods for provisioning Symmetrix VMAX storage:

- ◆ **Recommended:** This method relies on wizards to step you through the provisioning process, and is best suited for novice and advanced users who do not require a high level of customization (i.e., the ability to create their own volumes, storage groups, etc.).
- ◆ **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 you will have a masking view, in which the volumes in the storage group are masked to the initiators in the initiator group and mapped to the ports in the port group.



Before you begin:

- ◆ The Symmetrix VMAX system must already be configured.
- ◆ For instructions on provisioning Symmetrix DMX storage, refer to [Masking volumes](#) (page 281).

To provision Symmetrix VMAX storage:

- | Recommended | Advanced |
|--|---|
| <ol style="list-style-type: none"> 1. Create a host (initiator group)
Use the Create Host Wizard to group host initiators (HBAs). | <ol style="list-style-type: none"> 1. Create a host (initiator group)
Use the Create Host Wizard to group host initiators (HBAs) |

2. **Provision storage for the new host**
Use the Provision Storage Wizard, which will step you through the process of creating the storage group, port group, and masking view, and to optionally associate the storage group with a FAST policy.
2. **Create the volumes**
Create one or more volumes on the Symmetrix system.
3. **Create a storage group**
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 wizard's **Storage Group Type** to **Empty**, and then complete [Adding volumes to empty storage groups](#).
4. **Create a port group**
Group Fibre Channel and/or iSCSI front-end directors.
5. **Create a masking view**
Associate the initiator, storage, and port groups into a masking view.
6. **Associate the storage group with a FAST policy**
Optional: Associate the storage you created in step 3 with an existing FAST policy and assign a priority value for the association.

Using the Provisioning Storage wizard

This procedure explains how to provision Symmetrix VMAX storage using the Provisioning Storage wizard. The wizard steps you through the provisioning process, and is best suited for novice and advanced users who do not require a high level of customization (i.e., the ability to create their own volumes, storage groups, etc.). For users who want the ability to control every aspect of the provisioning process, refer to the **Advanced** procedure in See [Provisioning Symmetrix VMAX storage overview](#) on previous page..

Before you begin:

- ◆ The Symmetrix VMAX system must already be configured and you must already have a host (initiator group). For instructions on creating an initiator group, refer to See [Creating initiator groups \(Hosts\)](#) on page 290..
- ◆ For instructions on provisioning Symmetrix DMX storage, refer to [Masking volumes](#) (page 281).

To use the Provisioning Storage wizard:

1. Select **Provision Storage** from the **Common Tasks** pane to open the **Provision Storage** wizard.
2. Select a **Symmetrix** system and an existing **Host**.
3. Select whether to **Create a New Storage Group, Use an Existing Storage Group, or Expand an existing Storage Group**.
4. If using/expanding an existing storage group, select one and click **Next**.
5. Do the following, depending on operation you are performing:

Creating a New Storage Group:

- a. In the **Storage Group Name** field, you can either accept the system-generated name (host name with **_sg** appended to it) or type a new name over it. Storage group names must be unique from other storage groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (**_**), and dashes (**-**) are allowed. Storage group names are case-insensitive.
- b. Select whether to create a **Standard Storage Group** or a **Cascaded Storage Group**.
- c. If creating a Standard Storage Group, select the **Volume Type** to add to the storage group and click **Next**.

- d. Do the following, depending on the **Volume Type** you selected:

Virtual Volumes:

- a. Select the **Emulation** type for the volumes to add to the storage group.
- b. *Optional:* Select the **Thin Pool** containing the volumes to add to the storage group.
- c. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- d. *Optional:* Click **Show Advanced** and set any number of the following options:

- ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
- ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
 - **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.
- ◆ To **Preallocate capacity for each volume** you are adding to the storage group, select the option and specify the amount of each volume to preallocate, either **Full Volume** or **By Capacity**.
 - ◆ To **Persist preallocated capacity through reclaim or copy**, select the option.

- e. Click **Next**.

Regular Volumes:

- a. Select the **Disk Technology** on which the storage group will reside.
- b. Select the **Emulation** type for the volumes to add to the storage group.
- c. Select the **Protection** level for the volumes to add to the storage group.
- d. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- e. *Optional:* Click **Show Advanced** and select any number of the following options:

- ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
- ◆ To only use volumes from a specific disk group, select the **Disk Group**.
- ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new

Note volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
- **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

f. Click **Next**.

Template:

- a. Type or **Select** the name of the template.
- b. Type the **Number of Volumes** to add to the storage group.
- c. *Optional:* To name the volumes you are adding to the storage group, click **Show Advanced** and select one of the following **Volume Identifiers** and type a **Name:**

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- ◆ **Name Only** — All volumes will have the same name.
- ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

d. Click **Next**.

Cascaded Storage Group:

- a. Move one or more storage groups from the **Available Storage Groups** list to the **Selected Storage Groups** list.
- b. Click **Next**.

Using an Existing Storage Group:

- a. Select the **Storage Group** and click **Next**. If the storage group you selected is already in a masking view with a port group and an initiator group, you can skip to step 6.

- b. Select whether to use a **new port group** or an **existing port group**. If using a new port group, you can either accept the system-generated name (host name with `_pg` appended to it) or type a new name over it. Port group names must be unique from other port groups on the Symmetrix system and cannot exceed 64 characters. Only alphanumeric characters, underscores (`_`), and dashes (`-`) are allowed. Port group names are case-insensitive.
- c. *Optional*: Select **Include ports not visible to this host**.
- d. Select and clear ports in the Ports table.
- e. Click **Next**.
- f. In the **Masking View** field, you can either accept the system-generated name (host name with `_mv` appended to it) or type a new name over it. Masking view names must be unique from other masking views on the Symmetrix system and cannot exceed 64 characters. Only alphanumeric characters, underscores (`_`), and dashes (`-`) are allowed. Masking view names are case-insensitive.

Expanding an existing storage group:

- a. Select the **Storage Group** and click **Next**.
- b. Select whether to expand the group with Virtual Volumes, Regular Volumes, a Template, or by copying another volumes configuration.

- c. Do the following, depending on how you want to expand the group:

Virtual Volumes:

- a. Select the **Emulation** type for the volumes to add to the storage group.
- b. *Optional:* Select the **Thin Pool** containing the volumes to add to the storage group.
- c. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- d. *Optional:* Click **Show Advanced** and set any number of the following options:

- ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
- ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
- **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.
- ◆ To **Preallocate capacity for each volume** you are adding to the storage group, select the option and specify the amount of each volume to preallocate, either **Full Volume** or **By Capacity**.
- ◆ To **Persist preallocated capacity through reclaim or copy**, select the option.

- e. Click **Next**.

Regular Volumes:

- a. Select the **Disk Technology** on which the storage group will reside.
- b. Select the **Emulation** type for the volumes to add to the storage group.
- c. Select the **Protection** level for the volumes to add to the storage group.
- d. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- e. *Optional:* Click **Show Advanced** and select any number of the following options:

- ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
- ◆ To only use volumes from a specific disk group, select the **Disk Group**.
- ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new

Note volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
- **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

f. Click **Next**.

Template:

- Type or **Select** the name of the template.
- Type the **Number of Volumes** to add to the storage group.
- Optional:* To name the volumes you are adding to the storage group, click **Show Advanced** and select one of the following **Volume Identifiers** and type a **Name:**

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- ◆ **Name Only** — All volumes will have the same name.
- ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

Copy Volume Configuration:

- Select a volume to copy.
- Type the **Number of Volumes** to add to the group.
- Optional:* To name the volumes you are adding to the storage group, click **Show Advanced** and select one of the following **Volume Identifiers** and type a **Name:**

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- ◆ **Name Only** — All volumes will have the same name.

- ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.
- d. Click **Next**.
1. Verify your selections. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
 2. Click **Finish**.

Storage groups

Creating storage groups

This procedure explains how to create a storage group.

Before you begin:

- ◆ Storage groups require Engenuity 5874 or higher.
- ◆ The user must have Administrator or StorageAdmin permission.
- ◆ The maximum number of storage groups allowed on a Symmetrix system is 8,192.
- ◆ A storage group can contain up to 4,096 Symmetrix volumes.
- ◆ A Symmetrix volume can belong to more than one storage group.
- ◆ In this procedure you can optionally name the volumes you are adding the storage group. For more information, refer to [Setting volume names](#) on page 212

To create a storage group:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Click **Create** to open the **Create Storage Groups** wizard.
4. Type a **Storage Group Name**. Storage group names must be unique from other storage groups on the system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and hyphens (-) are allowed. Storage group names are case-insensitive.
5. Select a **Storage Group Type**. Possible values are:
 - ◆ **Standard Storage Group** — The default storage group type.
 - ◆ **Cascaded Storage Group** — Indicates this will be a parent storage group.
 - ◆ **Empty Storage Group** — Indicates the storage group will be empty. When you select this option the **Volume Type** field becomes inactive. Click **Finish**.
6. Do the following, depending on the type of storage group you are creating:

Standard Storage Group:

 - a. Select the type of volume (**Volume Type**) to include in the group. Possible values are:
 - ◆ **Virtual Volumes** — Creates the group using virtual volumes.
 - ◆ **Regular Volumes** — Creates the group using regular volumes. This option is not available for Symmetrix VMAX 10K/VMAXe systems.
 - ◆ **Template** — Creates the group using an existing storage template.

- b. Click **Next**.
- c. Do the following, depending on the volume type:

Virtual Volumes:

- a. Select the **Emulation** type for the volumes to add to the storage group.
- b. *Optional:* Select the **Thin Pool** to which the volumes will be added.
- c. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- d. *Optional:* Click **Show Advanced** and set any number of the following options:

- ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
- ◆ To **Preallocate capacity for each volume** you are adding to the storage group, select the option and specify the amount of each volume to preallocate, either **Full Volume** or **By Capacity**. If **By Capacity**, type a capacity value.
- ◆ To **Persist preallocated capacity through reclaim or copy**, select the option.
- ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
- **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

- e. Click **Next**.

Regular Volumes:

- a. Select the **Disk Technology** on which the storage group will reside.
- b. Select the **Emulation** type for the volumes to add to the storage group.
- c. Select the **Protection** level for the volumes to add to the storage group.
- d. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.

e. *Optional:* Click **Show Advanced** select any number of the following options:

- ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
- ◆ To only use volumes from a specific disk group, select the **Disk Group**.
- ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
- **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

f. Click **Next**.

Template:

- a. Type or select the **Template**.
- b. Type the **Number of Volumes** to add to the storage group.
- c. To name the volumes you are adding to the storage group, click **Show Advanced** and select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- ◆ **Name Only** — All volumes will have the same name.
- ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

d. Click **Next**.

Cascaded Storage Group:

- a. Click **Next**.
- b. Select one or more storage groups and click **Next**.

Empty Storage Group:

Click **Next**.

7. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
8. Click **Finish**.

Adding volumes to empty storage groups

This procedure explains how to add volumes to an empty storage groups.

Before you begin:

Storage groups require Enginuity 5874 or higher.

To add volumes to empty storage groups:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **View Details** to open its **Details** page.
4. In the **Related Objects** panel, click **Volumes - empty** to open the **Emptylist** view.
5. Click **Add** to open the **Add Volumes** dialog box.
6. Select the **Volumes** and click **Add**.

Expanding storage groups

Expanding a storage group refers to increasing the amount of storage in the group accessible to the mv cvvc b b asking view or in the FAST Policy.

Before you begin:

- ◆ Storage groups require Enginuity 5874 or higher.
- ◆ In this procedure you can optionally name the volumes you are adding the storage group. For more information, refer to [Setting volume names](#) on page 212

To expand a storage group:

1. Select the Symmetrix system.
2. Select **Storage** > **Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Expand** to open the **Expand Storage Group** wizard.
4. Select a method for expanding the storage group. Possible values are:
 - ◆ **Virtual Volumes** — Expands the group using virtual volumes.
 - ◆ **Regular Volumes** — Expands the group using regular volumes. This option is not available for Symmetrix VMAX 10K/VMAXe systems.
 - ◆ **Template** — Expands the group using a storage template.
 - ◆ **Copy Volume Configuration** — Expands the group by copying the configuration of volumes already in the group.
5. Click **Next**.
6. Do the following, depending on the method you are using:

Virtual Volumes:

- a. Select the **Emulation** type for the volumes to add to the storage group.
- b. *Optional:* Select the **Thin Pool** to which the volumes will be added.
- c. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- d. *Optional:* Click **Show Advanced** and set any of the following options:
 - ◆ To only use BCVs in the storage group, select **Use BCV volumes**.
 - ◆ To **Preallocate capacity for each volume** you are adding to the storage group, select the option and specify the amount of each volume to preallocate, either **By Percent** or **By Capacity**.
 - ◆ To **Persist preallocated capacity through reclaim or copy**, select the option.
 - ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when expanding storage groups with new volumes. Note that when expanding storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
 - **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.
- e. Click **Next**.

Regular Volumes:

- a. Select the **Disk Technology** on which the storage group will reside.
- b. Select the **Emulation** type for the volumes to add to the storage group.
- c. Select the **Protection** level for the volumes to add to the storage group.
- d. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
- e. *Optional:* Click **Show Advanced** and set any of the following options:
 - ◆ To only use BCVs in the storage group, select **BCV**.
 - ◆ To only use volumes from a specific disk group, select the **Disk Group**.
 - ◆ To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when expanding storage groups with new volumes. Note that when expanding storage groups with some new and

Note some existing volumes, the identifiers will only be applied to the new volumes.

- **Name Only** — All volumes will have the same name.
- **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

f. Click **Next**.

Template:

- a. Type or select the **Template**.
- b. Type the **Number of Volumes** to add to the storage group.
- c. To name the volumes you are adding to the storage group, click **Show Advanced** and select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when expanding storage groups with new volumes. Note that when expanding storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- ◆ **Name Only** — All volumes will have the same name.
 - ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.
- d. Click **Next**.

Copy Volume Configuration:

- a. Select a volume to copy.
- b. Type the **Number of Volumes** to add.
- c. To name the volumes you are adding to the storage group, click **Show Advanced** and select one of the following **Volume Identifiers** and type a **Name**:

Note This option is only available when creating storage groups with new volumes. Note that when creating storage groups with some new and some existing volumes, the identifiers will only be applied to the new volumes.

- ◆ **Name Only** — All volumes will have the same name.
- ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

- d. Click **Next**.
7. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
8. If the storage group you are expanding is part of one or more masking views, you can manually set the LUN values of the volumes you are adding or creating in the storage group by doing the following:
 - a. Click **Set LUNs**. Note that this option will not be available if the storage group is not part of a masking view.

The **LUN Value** column lists the LUN addresses of the volumes. The addresses of masked volumes are shown in normal text. The addresses of unmasked volumes are shown in italics, indicating that the addresses are suggestions.
 - b. To edit a suggested address, either double-click it or select it and click **Next Available**, which will increment the **Start LUN Address** to the next available address.
9. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
10. Click **Finish**.

Modifying storage groups

This procedure explains how to add/remove volumes from a storage group.

Before you begin:

- ◆ Storage groups require Enginuity 5874 or higher.
- ◆ The user must have Administrator or StorageAdmin permission.
- ◆ The maximum number of storage groups allowed on a Symmetrix system is 8,192.
- ◆ A storage group can contain up to 4,096 Symmetrix volumes.
- ◆ A Symmetrix volume can belong to more than one storage group.

To modify a storage group:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **View Details** to open its **Details** view.
4. Make changes in the **Properties** panel and click **Apply**.
5. In the **Related Object** panel, click **Volumes** to open the **Volumes** list view.
6. Do the following, depending on whether you are adding or removing volumes:

Adding volumes:

- a. Click **Add** to open the **Add Volumes to Storage Group** dialog box.
- b. Select the volumes to add to the storage group.
- c. *Optional:* If adding volumes to a storage group that is part of a masking view, click **Set Dynamic LUN Addresses** to open the **Set LUN Address - Storage Group** dialog box, in which you can manually assign the host LUN addresses for all the volumes you are adding to the group. When done, click **OK** to close the dialog box.
- d. Click **OK** in the **Add Volumes to Storage Group** dialog box.

Removing volumes:

- a. Select the volumes to remove from the storage group and click **Remove** to open the **Remove Volume** dialog box.
- a. *Optional:* If the volumes are mapped, select to **Unmap** them.
- a. *Optional:* If the volumes are thin, select to **Unbind** them.
- b. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Associating/Disassociating FAST policies and storage groups

The procedure for associating FAST policies and storage groups, depends on whether you are associating a storage group with a policy or policy with a storage group.

Note that you can only disassociate a storage group from a policy, not a policy from a storage group.

Before you begin:

- ◆ Storage groups and FAST policies can only be associated under the following conditions:
 - The storage group is not already associated with another policy. This restriction does not apply to Symmetrix systems running Enginuity 5876.
 - The target FAST policy needs to have a least one pool that is part of the source policy in re-association activity.
 - The volumes in the new storage group are not already in a storage group associated with a FAST policy.
 - The policy has at least one tier.
 - The storage group only contains meta heads; meta members are not allowed.
 - The storage group does not contain moveable volumes. When a storage group is associated with a policy, you cannot add non-moveable volumes to it. Non-moveable volumes include:
 - CKD EAV
 - DRV
 - SFS
 - iSeries, ICOS, ICL
 - SAVE volumes
 - VDEVs
 - Diskless volumes
- ◆ The storage group cannot contain a volume that is part of another storage group already associated with another policy.
- ◆ The Symmetrix system has less than the maximum number of allowed associations (8,192).

To associate a FAST policy with a storage group:

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** details view.
5. Select the policy and click **Associate Storage Group** to open the **Associate Storage Group** dialog box.
6. Select one or more storage groups from the **Available Storage Groups** list and click **Add** to move them to the **Associated Storage Groups** list, or click **Add All** to move all storage groups from the **Available Storage Groups** list to the **Associated Storage Groups** list.
7. Click either of the following:
 - ◆ **OK** to associate the storage group.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To have FAST factor the R1 volume statistics into move decisions made for the R2 volume, select **Enable FAST [VP/DP] RDF Coordination**. This attribute can be set on a storage group, even when there are no SRDF volumes in the storage group. This feature is only available if the Symmetrix system is part of an SRDF setup. Both R1 volumes and R2 volumes need to be running Enginuity version 5876 or higher for the FAST VP system to coordinate the moves. However, the setting of the RDF coordination attribute will not be prevented if one of the Symmetrix systems is running an Enginuity level lower than 5876.
2. Click **OK**.

To associate a storage group with a FAST policy:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Associate to FAST** to open the **Associate to FAST Policy** dialog box.
4. Select a policy and click **OK**.

To disassociate a storage group from a FAST Policy:

1. Select the Symmetrix system.
2. In the navigation bar, click **Storage** to open the **Storage** section.
3. Click **FAST** to open the **FAST** dashboard.
4. Select the **FAST Type**.
5. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** details view.
6. Select the policy and click **View Details** to open the policy's details view.
7. In the **Related Object** view panel, click **Storage Groups** to open the **Storage Groups for FAST Policy** details view.
8. Select the one or more storage groups and click **Disassociate**.
9. Click **OK** in the confirmation message.

Adding or removing storage groups for cascaded storage groups

This procedure explains how to add or remove child storage groups from parent storage groups.

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the parent storage group and click **View Details** to open its **Details** view.
4. In the **Related Object** panel, click **Child Storage Groups** to open the **Cascaded Storage Groups** list view.
5. Do the following, depending on whether you are adding or removing storage groups:
 - Adding storage groups:**
 - a. Click **Add Storage Groups** to open the **Add Storage Groups** dialog box.
 - b. Select one or more storage groups and click **OK**.
 - Removing storage groups:**
 - a. Select one or more storage groups and click **Remove**.
 - b. Click **OK** in the confirmation message.

Cloning storage groups

This procedure explains how to clone storage groups.

Before you begin:

- ◆ This feature requires the Engenuity 5876 Q2 SR or higher.
- ◆ The storage group must contain only thin volumes (except gatekeepers under 10 MB) and they must all be of the same type (either BCV or standard thin volumes (TDEVs)). This restriction also applies to cascaded storage groups, that is, all volumes in the parent and child storage groups must be thin and of the same type.
- ◆ The SYMAP_ALLOW_DEV_INT_MULTI_GRP option must be enabled. For instructions on enabling the option, refer to *Editing the Options file* in the *EMC Solutions Enabler Installation Guide*.
- ◆ Meta volumes are not supported.

To clone storage groups:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Clone** to open the **Clone Storage Group** dialog box.
4. Type the **Name of Device Group** to hold the target volumes.
5. Select the thin pool to which the target volumes will be bound (**Bind to Pool**). If the source storage group contains thin volumes bound to different thin pools, or if its a cascaded storage group with child storage groups containing volumes bound to different thin pools, selecting a single thin pool will result in all target volumes being bound to that single pool.
6. *Optional:* Clear the **Create Replica Storage Group** option in which case a storage group for the target volumes will not be created. Leaving the option selected allows you to optionally change the name of replica storage group. Changing the name will also change the target volume storage group name.
7. *z/OS Only:* If the storage group contains CKD volumes, type a **New SSID** for the target, or click **Select ...** to open a dialog from which you can select an SSID
8. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. Select mode in which to create the clone session **Clone Copy Mode**. The mode you specify here will override the mode specified in the system preferences. Possible values are:
 - ◆ **Copy No Diff** — Creates a nondifferential (full) copy session in the background.
 - ◆ **PreCopy No Diff** — Creates a nondifferential (full) copy session in the background before the activate starts.
 - ◆ **Copy Diff** — Creates a differential copy session in the background. In differential copy sessions, only those volume tracks that have changed since the full clone was performed are copied (that is, only new writes to the source volume will be copied).
 - ◆ **PreCopy Diff** — Creates a differential copy session in the background before the activate starts. In differential copy sessions, only those volume tracks that have changed since the full clone was performed are copied (that is, only new writes to the source volume will be copied).
 - ◆ **VSE No Diff** — Creates a VP Snap Copy session.
2. Select a method for creating the clones (**Clone Creation**). Possible values are:
 - ◆ **Best Effort** — Specifies to make every effort to use existing volumes as the targets, and then create new volumes as necessary.
 - ◆ **Find Existing** — Specifies to only use exiting volumes as the targets.

3. Select the type of volumes to use as the targets (**Clone Targets**).
4. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Deleting storage groups

This procedure explains how to delete storage groups.

Before you begin:

- ◆ You cannot delete a storage group that is part of a masking view or associated with a FAST Policy.
- ◆ Storage groups require Enginuity 5874 or higher.

To delete a storage group:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Delete**.
4. Click **Delete** in the confirmation message.

Setting host I/O limits

Host I/O limits is a feature that can be used to limit the amount of Front End (FE) Bandwidth and I/Os per second (IOPs) that can be consumed by a set of Symmetrix volumes over a set of director ports. The bandwidth and I/Os against the set of volumes over the set of director ports will be monitored by the Symmetrix system to ensure that it will not exceed the user specified maximum bandwidth or maximum IOPs placed on these. This feature allows you to place limits on the FE Bandwidth and IOPs consumed by applications on the Symmetrix system.

Host I/O limits are defined as storage group attributes – the maximum bandwidth (in MB per second) and the maximum IOPs (in I/Os per second). For a cascaded storage group, a Host I/O limit can be added for either the parent or the child storage group, but not for both. If a parent storage group has a limit, the limit will be shared among all its child storage groups.

The Host I/O limit for a storage group can be either active or inactive, only the active Host I/O limit can limit the FE bandwidth and IOPs of the volumes in a storage group. The Host I/O limit will become active when a provisioning view is created using the storage group and will become inactive when the view is deleted. When a view is created on a parent storage group with a Host I/O limit, the limit will be shared among all the volumes in all child storage groups. Note that when a view is created on a parent storage group, even if no limit is defined on the parent storage group, you will not be able to define a limit for any of the child storage groups.

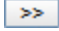
The Host I/O limit of the storage group will apply to all the director ports of the port group in the provisioning view. The Host I/O limit is divided equally among all the directors in the port group independent of the number of ports on each director. For this reason it is recommended that you configure only one of the ports of a director in the port group.

Before you begin:

The Symmetrix system must be running Enginuity 5876 Q4 2012 SR or higher.

For more information on setting host I/O limits, refer to the *EMC Solutions Enabler Symmetrix Array Management CLI Product Guide*.

To set host I/O limits:

1. Select a Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group, click more , and select **Set Host I/O Limits** to open **Set Host I/O Limits** dialog box.
4. Select and type values for one or both of the following:
 - ◆ **Set Host I/O Limit (MB/Sec)** — Maximum bandwidth (in MB per second). Valid values range from 1 MB/sec to 100,000 MB/sec.
 - ◆ **Set Host I/O Limit (IO/Sec)** — Maximum IOPs (in I/Os per second). Valid values range from 100 IO/Sec to 100,000 IO/sec, in 100 increments.
5. To configure a dynamic distribution of host I/O limits, **Set Dynamic Distribution** to one of the following; otherwise, leave this field set to **Never** (default). This feature requires the Enginuity 5876 Q2 2013 SR.
 - ◆ **Always** — Enables full dynamic distribution mode. When enabled, the configured host I/O limits will be dynamically distributed across the configured ports, thereby allowing the limits on each individual port to adjust to fluctuating demand.
 - ◆ **Failure** — Enables port failure capability. When enabled, the fraction of configured host I/O limits available to a configured port will adjust based on the number of ports currently online.
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now

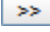
Managing VP compression on storage groups

The following explains how to manage VP compression on the thin volumes in a storage group.

Before you begin:

This feature requires Enginuity 5876 Q42012 SR.

To manage VP compression on storage groups:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group, click , and select **VP Compression** to open the **VP Compression** dialog box.
4. Select one of the following compression actions:
 - ◆ **Start Compression** — Starts compressing the thin volumes in the storage group.
 - ◆ **Stop Compression** — Stops compressing the thin volumes in the storage group.
 - ◆ **Start Uncompression** — Starts uncompressing the thin volumes in the storage group.
 - ◆ **Stop Uncompression** — Stops uncompressing the thin volumes in the storage group.









Viewing storage groups

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.

The **Storage Groups** list view allows you to view and manage storage groups on a Symmetrix system.

Note There are multiple ways to open this view. Depending on the one you used, some of the following properties and controls may not appear.

The following properties display:

- ◆ **Name** — Name of the storage group, accompanied by an icon indicating the storage group type. Possible types are:
 -  Standalone storage group
 -  Child storage group
 -  Parent storage group
 -  Storage group in masking view
 -  Child storage group in masking view
 -  Parent storage group in masking view
- ◆ **Parent** — Indicates if this storage group is a parent in a cascaded relationship ( = yes).
- ◆ **Child** — Indicates if this storage group is a child in a cascaded relationship ( = yes).
- ◆ **Child SGs** — Number of child storage groups contained in this (parent) storage group. This field will display 0 for child storage groups.
- ◆ **FAST Policy** — Policy associated with the storage group.
- ◆ **Capacity** — Total capacity of the storage group in GB.
- ◆ **Volumes** — Number of volumes contained in the storage group.
- ◆ **Masking Views** — Number of masking views associated with the storage group.

The following controls are available:

- ◆ **Create** — See [Creating storage groups](#) on page 148.
- ◆ **Expand** — See [Expanding storage groups](#) on page 152.
- ◆ **View Details** — See [Viewing storage group details](#) on next page.
- ◆ **Delete** — See [Deleting storage groups](#) on page 162.
- ◆ **Associate to FAST** — See [Associating/Disassociating FAST policies and storage groups](#) on page 183.
- ◆ **Reassociate to FAST** — See [Reassociating FAST polices and storage groups](#) on page 186.
- ◆ **Dissociate** — See [Associating/Disassociating FAST policies and storage groups](#) on page 183.
- ◆ **Replication QoS** — See [Setting replication QoS](#) on page 137.
- ◆ **Assign Symmetrix Priority** — See [Assigning Symmetrix priority](#) on page 205. (Does not apply/display for Symmetrix VMAX 10K (VMAXe) systems)
- ◆ **Tag for RecoverPoint** — See [Tagging and untagging volumes for RecoverPoint](#) on page 410.

- ◆ **Untag for RecoverPoint** — See [Tagging and untagging volumes for RecoverPoint](#) on page 410.
- ◆ **Bind SG** — See [Binding/Unbinding/Rebinding thin volumes](#) on page 244.
- ◆ **Unbind SG** — See [Binding/Unbinding/Rebinding thin volumes](#) on page 244.
- ◆ **Rebind SG** — See [Binding/Unbinding/Rebinding thin volumes](#) on page 244.
- ◆ **Pin SG** — See [Pinning and unpinning volumes](#) on page 191.
- ◆ **Unpin SG** — See [Pinning and unpinning volumes](#) on page 191.
- ◆ **VLUN Migration** — See [Migrating regular storage group volumes](#) on page 384.
- ◆ **Assign Dynamic Cache Partition** — See [Assigning volumes to dynamic cache partitions](#) on page 131.
- ◆ **VP Compression** — See [Managing VP compression on storage groups](#) on page 163.
- ◆ **Set Host I/O Limits** — [Setting host I/O limits](#) on page 162

Viewing storage group details

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **View Details** to open its **Details** view.

The storage group **Details** view allows you to view and manage a storage group. It contains [Properties](#), [Related Objects](#), [Performance Views](#), and [FAST Compliance Report](#) panels.

Note There are multiple ways to open this view. Depending on the one you used, some of the following properties and controls may not appear.

Properties panel

The following properties display:

- ◆ **Name** — Name of the storage group. To rename the storage group, type a new name over the existing and click **Apply**. Storage group names must be unique from other storage groups on the system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Storage group names are case-insensitive.
- ◆ **FAST Policy** — Policy associated with the storage group. To associate the storage group with a different policy, select a new policy and click **Apply**.
- ◆ **FAST Priority** — The priority value currently assigned to the storage group within the tier. Storage groups associated with a policy are assigned a priority value that determines the order in which the tier will service them during conflicts. Possible values range from 1 (the highest) to 3 (the lowest). To change the priority, select a new priority and click **Apply**. This field only displays when the storage group is associated with a FAST Policy.
- ◆ **Total Capacity** — Total capacity of the storage group in GB.
- ◆ **Host Name** — Hosts associated with the storage group.

- ◆ **Volumes** — Number of volumes in the storage group.
- ◆ **Masking Views** — Number of masking views associated with the storage group.
- ◆ **Enable FAST RDF Coordination** — Enables or disables SRDF coordination. When enabled, FAST factors the R1 volume statistics into the move decisions that are made on the R2 volume. This attribute can be set on a storage group, even when there are no SRDF volumes in the storage group. To enable or disable FAST RDF Coordination, select or clear the option and click **Apply**. This feature is only available if the Symmetrix system is part of an SRDF setup. Both R1 devices and R2 devices need to be running Enginuity version 5876 or higher for the FAST VP system to coordinate the moves. However, the setting of the RDF coordination attribute will not be prevented if one of the Symmetrix systems is running an Enginuity level lower than 5876. This field only displays when the storage group is associated with a FAST Policy.
- ◆ **Last Updated** — Timestamp of the most recent changes to the storage group.
- ◆ **Child Storage Groups** — Number of child storage groups contained in this (parent) storage group. This field only displays for parent storage groups.
- ◆ **Parent Storage Groups** — Number of storage groups of which this storage group is a child. This field only displays for child storage groups.
- ◆ **Host I/O Limit** — Whether the host I/O limit feature is enabled. For more information, see [Setting host I/O limits](#) on page 162.
- ◆ **Set Host I/O Limit (MB/Sec)** — Enables you to specify the maximum bandwidth.
- ◆ **Host I/O Limit (MB/Sec)** — Maximum bandwidth (in MB per second). Valid values range from 1 MB/sec to 100,000 MB/sec.
- ◆ **Set Host I/O Limit (IO/Sec)** — Enables you to specify the maximum IOPs.
- ◆ **Host I/O Limit (IO/Sec)** — Maximum IOPs (in I/Os per second). Valid values range from 100 IO/Sec to 100,000 IO/sec.
- ◆ **Set Dynamic Distribution** — Enables/disables dynamic distribution of host I/O limits. This feature requires the Enginuity 5876 Q2 2013 SR. Possible values are:
 - ◆ **Always** — Enables full dynamic distribution mode. When enabled, the configured host I/O limits will be dynamically distributed across the configured ports, thereby allowing the limits on each individual port to adjust to fluctuating demand.
 - ◆ **Failure** — Enables port failure capability. When enabled, the fraction of configured host I/O limits available to a configured port will adjust based on the number of ports currently online.
 - ◆ **Never** — Disables this feature (Default).

The following controls are available:

- ◆ **Create** — See [Creating storage groups](#) on page 148.
- ◆ **Expand** — See [Expanding storage groups](#) on page 152.
- ◆ **Delete** — See [Deleting storage groups](#) on page 162.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the storage group.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Objects panel

The **Related Objects** panel links you to views displaying objects contained in and associated with the storage group. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Masking Views - 2** will open a view listing the two masking views associated with the storage group.

Performance Views panel

The **Performance** panel links you to the performance monitor and analyze views for the storage group.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

FAST Compliance Report panel

The **FAST Compliance Report** provides the following information:

- ◆ **Tier** — Tier name.
- ◆ **Protection** — RAID protection level.
- ◆ **Technology** — Disk technology.
- ◆ **Max SG Demand (%)** — The percentages for the storage group per tier as defined in the FAST Policy.
- ◆ **Limit (GB)** — Upper limit in GB per tier for the storage group.
- ◆ **Fast SG Used (GB)** — Current occupancy in GB of the storage group in the tier.
- ◆ **Growth (GB)** — Per the FAST policy, how much more the storage group can grow in GB on a given tier. This also indicates compliance. If the growth is negative that means the storage group has exceeded the capacity limit for this tier and is out of compliance.

Fully Automated Storage Tiering

Understanding FAST

Fully Automated Storage Tiering (FAST) is Symmetrix software that runs background algorithms to continuously analyze the utilization (busy rate) of the Symmetrix system volumes. The FAST controller processes the algorithm data, and generates plans for moving and swapping data volumes to fine tune performance and reduce costs. FAST can move the most-used data to the fastest (and most expensive) storage, such as Enterprise Flash Drives (EFD), the least-used data to the slowest (and least expensive) storage, such as SATA drives, while maintaining the remaining data on Fibre Channel (FC) drives, based on user-defined Symmetrix tiers and FAST policies. The objective of tiered storage is to minimize the cost of storage, while improving or maintaining performance, by putting the right data, on the right Symmetrix tier, at the right time.

Enginuity 5876 introduced Federated Tiered Storage (FTS), which allows the virtualization of external storage as an external disk (eDisk). Adding the eDisk to the Symmetrix system makes its capacity available to the Symmetrix system as an external spindle. FAST VP supports tiers of externally provisioned VP pools. Encapsulated volumes are not supported. There is no support for externally provisioned or encapsulated (standard) volumes with FAST. The new order for the fastest to the slowest tiers is EFD, FC, SATA, and external tiers.

Enginuity 5876 Q4 2012 SR supports associating FTS tiers with a technology type. The technology associated with the FTS tier indicates the tier's expected performance to the FAST VP controller. This enables you to place the FTS tier at the right location for the expected performance of the external tier.

After configuration, FAST can be set to move/swap data automatically or with user approval. All three drive technologies (EFD, FC, SATA), or external disks (eDisks) are not required in the Symmetrix system to use FAST.

For information on FAST concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

FAST versions

There are two versions of FAST: FAST and FAST for Virtual Pools (FAST VP). The following table identifies the differences between the versions:

Table 4: Comparison of FAST and Fast Virtual Pools (VP)

FAST	FAST Virtual Pools
Requires Enginuity 5874 or higher	Requires Enginuity 5875 or higher
Supports standard volumes	Supports thin volumes
Supports FBA and CKD volume emulation	Supports FBA volume emulation Enginuity 5876 supports thin CKD 3390 and thin IBM i 512 volumes
Disk group provisioning (DP) tiers: contain disk groups	Virtual pool (VP) tiers: contain thin pools

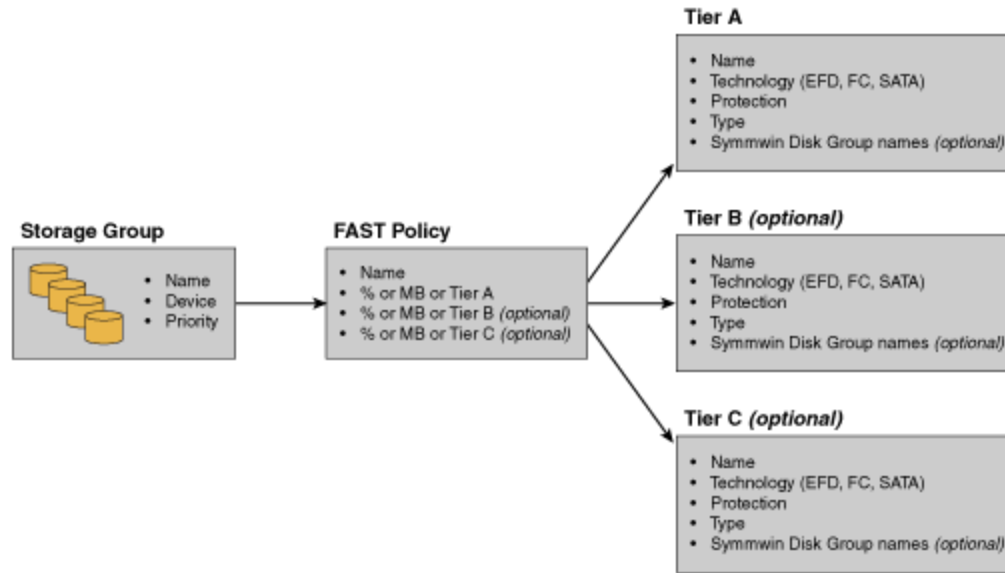
FAST	FAST Virtual Pools
DP modes: Auto Approve and User Approve	VP modes: Auto Approve or None
User visible data movement plans and history	No plans or history generated
Federated Tiered Storage (eDisks) not supported	Supports Federated Tiered Storage (eDisks) with Enginuity 5876 or higher
Thin volume/thin pool compression not supported	Supports compression for thin volumes and thin (VP) pools
Up to three tiers per policy supported	Up to four VP tiers per policy with Enginuity the 5876 Q4 2012 SR

Configuration overview

The following are the basic steps for configuring FAST in a Symmetrix system:

1. **Create a set of Symmetrix tiers.** A Symmetrix tier is a specification of the type of storage (EFD, FC, SATA, or eDisks) and resources (disk groups/virtual pools) from which storage will be selected. With FAST, from one to four tiers are grouped together into a policy.
2. **Create a set of storage groups.** A storage group is a group of volumes. When used with FAST, a storage group is associated with a FAST policy and assigned a priority.
3. **Create a FAST policy.** A FAST policy is a set 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 how much data from a storage group associated with the policy is allowed to reside on the tier.
4. **Associate the storage groups with the FAST policy.** A storage group can only be associated with one policy; however, one policy can be associated with multiple storage groups.

The association between a storage group, FAST policy, and Symmetrix tiers is illustrated in the following figure:



GEN-001 191

Managing FAST

Time windows define when the FAST controller should collect performance information or execute data movement. Time windows are described in [Understanding time windows](#).


Control parameters define the numbers of and types of volumes, the modes of operation, the thresholds for data movement, and the analysis time period. Control parameters are described in [FASTcontroller](#) (page 189).

Monitoring FAST

Once you have configured FAST on a Symmetrix system, you can use the FAST dashboard to monitor its status, as described in [Monitoring FAST](#) (page 172).

Monitoring FAST

The **FAST** dashboard provides you with a single place from which to examine FAST data, including status and demand reports, and information on policies and storage groups under FAST control. In addition, you can manage the FAST polices on a Symmetrix system.

Note To refresh the information displayed in the **FAST** dashboard, click refresh Symmetrix () in the title bar, or exit/return to the **FAST** dashboard.

To access the FAST dashboard:

1. Select the Symmetrix system.
2. Select **Storage > FAST**.

FAST dashboard

The **FAST** dashboard contains the following components:

Fast Type

If the Symmetrix system is licensed for both FAST and FAST VP, use this option to toggle the monitor between the two features; otherwise, the licensed feature will appear selected for you. The **FAST Status Report**, **FAST Policies**, **Tiers Demand Report**, and **Storage Groups Under FAST Control** view panels will update according to the feature you select.

FAST Status Report

Displays information on the current state of FAST and some of the more important FAST settings, including:

- ◆ **State** — The current state of FAST on the Symmetrix system. Possible values are:
 - **Enabled** — The FAST controller is Enabled.
 - **Disabled** — The FAST controller is Disabled.
 - **Disabling** — The FAST controller is transitioning from Enabled to Disabled.
 - **Enabling** — The FAST controller is transitioning from Disabled to Enabled.
 - **Disabled with Error** — The FAST controller is disabled with an error.
 - **Degraded** — The FAST controller is activated but not fully functional because of any of the following reasons: missing DRV volumes or because of some other reason. When degraded, moves are possible; however, static swaps are not.
- ◆ **Data Movement Mode** — The mode in which the FAST controller is operating. Possible values are:
 - **Automatic** — The FAST controller will continuously perform data movement within the time window, without user intervention.
 - **User Approval** — The FAST controller will generate plans, but not perform any movements without user approval. This value only applies to FAST DP.
 - **Off** — The FAST controller will not perform any data movements for thin volumes. This value only applies to FAST VP.
- ◆ **Current Activities** — The current FAST activity. Possible values are:
 - **RunningPlan** — FAST is currently moving or swapping data according to plan.

- **Idle** — FAST is enabled, but not running a plan.
- ◆ **Performance Time Window** — Indicates whether FAST is operating within a defined performance time window. Possible values are:
 - **Closed** — Indicates that FAST is operating within a performance time window.
 - **Open** — Indicates that FAST is operating outside a performance time window (that is, the last performance time window has expired, the next performance time window has yet to start, or there are no performance time windows defined).
- ◆ **Move Time Window** — Indicates whether FAST is operating within a defined move time window. Possible values are:
 - **Closed (exclusion time window)** — Indicates that FAST is operating within a move time window.
 - **Open (inclusion time window)** — Indicates that FAST is operating outside a move time window (that is, the last move time window has expired, the next move time window has yet to start, or there are no move time windows defined).

FAST Policies

Allows you view and manage the FAST policies on the Symmetrix system. This view panel includes the following attributes:

- ◆ **Manage Policies** — Opens the **FAST Policies** list view, from which you can manage the FAST Policies on the Symmetrix system.
- ◆ **Policy List** — Lists the policies on the Symmetrix system, including the following:
 - **Policy Name** — Name of the policy.
 - **Tier 1** — Symmetrix tier associated with the policy.
 - **Tier 1 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 1.
 - **Tier 2** — Symmetrix tier associated with the policy.
 - **Tier 2 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 2.
 - **Tier 3** — Symmetrix tier associated with the policy.
 - **Tier 3 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 3.
 - **Tier 4** — Symmetrix tier associated with the policy.
 - **Tier 4 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 4.

This view panel displays up to 4 tiers is for FAST VP policies and up to 3 tiers for FAST policies.

- **Associated Storage Group** — Storage group associated with the policy.

Tiers Demand Report

Allows you view the demand on each tier in the Symmetrix system, as either a chart or a table.

The chart format provides the following information:

- ◆ **Used** — The amount of storage that has already been used on the tier, in GB.
- ◆ **Free** — The amount of free/unused storage on the tier, in GB.

- ◆ **Max SG Demand** — The maximum amount of storage that FAST could put on the tier, in GB.
- ◆ **Available** — The amount of storage available to FAST on the tier, in GB.

The table format provides the same information as the chart format, with the following additions:

- ◆ **Tech + Port** — The type of disk on which the tier resides and the tier's RAID protection level.
- ◆ **Excess** — The amount of excess storage available if FAST reaches the maximum of all the percentages in the policies associated with storage groups on the tier, in GB. This value can be either positive or negative. A positive value indicates that there will be enough excess storage available for all the storage groups. A negative value indicates that there will not be enough excess storage available and storage groups will be competing against one another.

Storage Groups Under FAST Control

Allows you to view information on all of the storage groups under FAST control on the Symmetrix system, including:

- ◆ **Storage Group Name** — The name of the storage group.
- ◆ **FAST Policy** — The FAST Policy associated with the storage group.
- ◆ **Capacity Break Down Per Tier** — A graphic representation of where the storage group's capacity resides in relation to the tiers in the FAST Policy.
- ◆ **Compliant** — Indicates whether the storage group is within compliance. A storage group is compliant if all its volumes exist only on the tiers defined in the policy and the percentage capacity of all tiers it occupies are within the upper limits of the tier capacities specified in the policy. **Yes** indicates compliance. **No** indicates non-compliance.

Symmetrix tiers

Creating Symmetrix tiers

This procedure explains how to create Symmetrix tiers.

Before you begin:

- ◆ The maximum number of tiers that can be defined on a Symmetrix 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 will:
 - Lead to a mix of static and dynamic tier definitions in the same technology.
 - Partially overlap with an existing tier. Two tiers partially overlap when they share only a subset of disk groups. For example, TierA partially overlaps with TierB when TierA contains disk groups 1 & 2 and TierB contains only disk group 2. (Creating TierA will fail.)

To create a Symmetrix tier:

1. Select the Symmetrix system.
2. Select **Storage > Tiers** to open the **Tiers** list view.

3. Click **Create** to open the **Create Tier** dialog box. When this dialog box first opens, the chart displays the configured and unconfigured space on the selected Symmetrix system. Once you select a disk group or thin pool, this chart will display the configured and unconfigured space of the selected object.
4. Type a **Tier Name**. Tier names must be unique and cannot exceed 32 characters. Only alphanumeric characters, hyphens (-), and underscores (_) are allowed, however, the name cannot start with a hyphen or an underscore. Each tier name must be unique per Symmetrix system (across both DP and VP tier types), ignoring differences in case.
5. If the Symmetrix system on which you are creating the tier is licensed to perform FAST and FAST VP operations, select a **Tier Type**. Possible values are:
 - ◆ **DP Tier** — A disk group tier is a set of disk groups with the same technology type. A disk group tier has a disk technology type and a protection type. To add a disk group to a tier, the group must only contain volumes on the tier's disk technology type and match the tier protection type.
 - ◆ **VP Tier** — A virtual pool tier is a set of thin pools. A virtual pool tier has a disk technology type and a protection type. To add a thin pool to a tier, the thin pool must only contain DATA volumes on the tier's disk technology type and match the tier protection type.
6. If creating a VP tier, select the **Emulation** type of the thin pools to include in the tier. Only thin pools containing volumes of this emulation type will be eligible for inclusion in the tier.
7. Select the type of **Disk Technology** on which the tier will reside. Only disk groups or thin pools on this disk technology will be eligible for inclusion in the tier.
8. If you selected **External** disk technology for the tier, then select the type of **External Technology**.
9. Select the RAID **Protection** level for the tier. Only disk groups or thin pools on this disk technology will be eligible for inclusion in the tier.
10. Depending on the type of tier you are creating, select the disk groups or virtual pools to include in the tier.
11. If creating a DP tier, you can optionally specify to automatically add all future disk groups on matching disk technology to this tier. To do this, click **Show Advanced**, and select the option. Tiers created in this manner are considered dynamic tiers. Tiers created without this option are considered static tiers.
12. Click **OK**.

Modifying Symmetrix tiers

This procedure explains how to add and remove disk groups/thin pools from Symmetrix tiers.

Before you begin:

- ◆ You can only modify tiers that are not part of a policy. For instructions on removing a tier from a policy, refer to *Modifying FAST policies*.
- ◆ You cannot create blank tiers in Unisphere for VMAX (that is, tiers without disk groups or thin pools); however, you can use Unisphere to add disk groups or thin pools to blank tiers that were created in Solutions Enabler.

To modify a Symmetrix tier:

1. Select the Symmetrix system.
2. Select **Storage > Tiers** to open the **Tiers** list view.
3. Select the tier and click **Edit** to open the **Edit Tier** dialog box.
4. Add or remove disk groups|thin pools by selecting/clearing the corresponding check box.
5. If you are adding disk groups to a partially overlapped disk group tier (that is, a disk group tier that shares a subset of its disk groups with other disk group tiers) you must use the propagate option. To do this, click **Show Advanced**, and select **Adjust all disk group tiers**.
6. Click **OK**.

Renaming Symmetrix tiers

This procedure explains how to rename Symmetrix tiers.

Before you begin:

Tier names must be unique and cannot exceed 32 characters. Only alphanumeric characters, hyphens (-), and underscores (_) are allowed, however, the name cannot start with a hyphen or an underscore. Each tier name must be unique per Symmetrix system (across both DP and VP tier types), ignoring differences in case.

To rename a Symmetrix tier:

1. Select the Symmetrix system.
2. Select **Storage > Tiers** to open the **Tiers** list view.
3. Select the tier and click **View Details** to open its **Details** view.
4. Type a new name for the tier.
5. Click **Apply**.

Deleting Symmetrix tiers

This procedure explains how to delete Symmetrix tiers.

Before you begin:

You cannot delete tiers that are already part of a policy. To delete such a tier, you must first remove the tier from the policy. For instructions, refer to *Modifying FAST policies*.

To delete a Symmetrix tier:

1. Select the Symmetrix system.
2. Select **Storage > Tiers** to open the **Tiers** list view.
3. Select the tier and click **Delete**.
4. Click **OK** in the confirmation message.

Viewing Symmetrix tiers

1. Select the Symmetrix system.
2. Select **Storage > Tiers** to open the **Tiers** list view.

The **Tiers** list view allows you to view and manage the tiers on a Symmetrix system.

The following properties display:

- ◆ **Name** — Name of the tier.
- ◆ **Type** — Tier type. Possible values are:
 - **DP - Disk Group Provisioning** — A disk group tier is a set of disk groups with the same technology type. A disk group tier has a disk technology type and a protection type. To add a disk group to a tier, the group must only contain volumes on the tier's disk technology type and match the tier protection type.
 - **VP Virtual Provisioning** — A virtual pool tier is a set of thin pools. A virtual pool tier has a disk technology type and a protection type. To add a thin pool to a tier, the thin pool must only contain DATA volumes on the tier's disk technology type and match the tier protection type.
- ◆ **Technology** — Disk technology on which the tier resides.
- ◆ **Disk Location** — Indicates whether the tier is internal or external (eDisk).
- ◆ **Emulation** — Emulation type of the thin pools in the tier.
- ◆ **Used Capacity** — Amount of storage that has already been used on the tier, in GB.
- ◆ **Capacity** — Amount of free/unused storage on the tier, in GB.
- ◆ **Protection** — RAID protection level assigned to the volumes in the tier.

The following controls are available:

- ◆ **Create** —See [Symmetrix tiers](#) on page 174.
- ◆ **Edit** —See [Modifying Symmetrix tiers](#) on page 175.
- ◆ **View Details** —See [Viewing Symmetrix tier details](#) below.
- ◆ **Delete** —See [Deleting Symmetrix tiers](#) on page 176.

Viewing Symmetrix tier details

1. Select the Symmetrix system.
2. Select **Storage > Tiers** to open the **Tiers** list view.
3. Select the tier and click **View Details** to open its **Details** view.

The tier **Details** view allows you to view and manage a Symmetrix tier. It contains [Properties](#) below, [Related Objects](#) on the facing page, [Performance](#), and [Graphs](#) on page 180 panels. There are multiple ways to open this view. Depending on the method used, not all of the following may apply.

Properties

The following properties display:

- ◆ **Name** — Name of the tier.

[OutOfTier]: If on a given technology there exists volumes that do not reside on any tier they will be shown as [OutOfTier]. This can happen when the protection type of volumes does not match the tier protection type, or when tiers are only defined on a subset of disk groups in a technology.
- ◆ **Is Static** — Whether the tier is static (Yes) or dynamic (No). With a dynamic tier, the FAST controller will automatically add all future disk groups on matching disk technology to the tier. Tiers without this option enabled are considered static.
- ◆ **Type** — Tier type. Possible values are:
 - **DP** — A disk group tier is a set of disk groups with the same technology type. A disk group tier has a disk technology type and a protection type. To add a disk group to a tier, the group must only contain volumes on the tier's disk technology type and match the tier protection type.
 - **VP** — A virtual pool tier is a set of thin pools. A virtual pool tier has a disk technology type and a protection type. To add a thin pool to a tier, the thin pool must only contain DATA volumes on the tier's disk technology type and match the tier protection type.
- ◆ **Technology** — Disk technology on which the tier resides.

If **Disk Location** is **External** then **External Technology** property displays and the technology value can be changed.
- ◆ **RAID Protection** — RAID protection level assigned to the volumes in the tier.

- ◆ **Attribute** — Status of the tier on the technology type. Possible values are:
 - Tier in a FAST Policy associated with storage groups.
 - Tier in a FAST Policy unassociated with storage groups.
 - Tier not in any FAST Policy.
- ◆ **Total Capacity** — Amount of free/unused storage on the tier, in GB.
- ◆ **Free Capacity** — Unconfigured space in Gigabytes in this tier. Free capacity for each disk group in the tier will only count toward tier free capacity if the disk group has enough usable disks to support the tier target protection type.
- ◆ **FAST Usage** — Sum of hypers of all volumes in FAST storage group with matching RAID protection that reside on this tier.
- ◆ **FAST Free** — If the tier is in a FAST policy associated with a storage group, the FAST Free capacity in Gigabytes is the sum of FAST Usage, Free capacity and Space occupied by Not Visible Devices (Unmapped/Unmasked).
If the tier is not in any FAST policy or in policies where none of the policies are associated to a storage group, then the FAST Available capacity is same as FAST Usage.
- ◆ **Maximum SG Demand** — The calculated upper limit for the storage group on the tier.
- ◆ **Excess** — Difference between FAST Free and Max SG Demand. If the tier is not in a FAST policy or in policies where none of the policies are associated to a storage group, then this value is Not applicable.
- ◆ **Dynamic Disk Location Capable** — Whether the tier allows for dynamic disk location.
- ◆ **VP Compression Capable** — Whether the tier allows for VP compression.

The following controls are available:

- ◆ **Create** — See [Symmetrix tiers](#) on page 174.
- ◆ **Edit** — See [Modifying Symmetrix tiers](#) on page 175.
- ◆ **Delete** — See [Deleting Symmetrix tiers](#) on page 176.
- ◆ **Apply** — Applies changes made to the tier name.
- ◆ **Cancel** — Cancels changes made to the tier name.

Related Objects

The **Related Objects** panel links you to views displaying objects contained in tier. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Disk Groups- 3** opens a view listing the three disk groups in the policy.

Performance panel

The **Performance** panel links you to the performance monitor and analyze views for the tier. This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Graphs

The **Graphs** panel provides a graphic representation of the tier's used capacity over free space.

FAST policies

Creating FAST policies

A FAST policy is a set 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 how much data from a storage group associated with the policy is allowed to reside on the tier.

Storage groups are sets of volumes. Storage groups define the volumes used by specific applications. Storage groups are associated with FAST policies, and all of the volumes in the storage group come under FAST control. The FAST controller can move these volumes (or data from the volumes) between tiers in the associated policy.

A storage group associated with a FAST policy may contain standard volumes and thin volumes, but the FAST controller will only act on the volumes that match the type of tier contained in the associated policy. For example, if the policy contains thin tiers, then the FAST controller will only act on the thin volumes in the associated storage group.

Before you begin:

- ◆ The maximum number of policies allowed per Symmetrix system is 256.
- ◆ Policies must contain either disk group tiers or virtual pool tiers, but not a combination of both disk group and virtual pool tiers.
- ◆ Disk group tier policies can contain from one to three tiers.
- ◆ Virtual pool tier policies can contain from one to four tiers. Only one out of the four tiers can be an external tier.
- ◆ Each tier must be unique and there can be no overlapping disk groups/thin pools.
- ◆ The first tier added to a policy determines the type of tier the policy will contain.
- ◆ A policy cannot have an empty tier.
- ◆ You cannot create blank policies (that is, policies without at least one tier) in Unisphere for VMAX; however, you can create such policies in Solutions Enabler. The *Solutions Enabler Symmetrix Array Controls CLI Product Guide* contains instructions on creating blank policies. Unisphere does allow you to manage blank policies.
- ◆ You cannot add a standard tier to a policy if it will result in a configuration where two tiers share a common disk group.

To create a FAST Policy:

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.

3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type** to which the policy will apply.
4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** details view.
5. Click **Create** to open the **Create FAST Policy** dialog box.
6. Type a **Policy Name**. Policy names must be unique and cannot exceed 32 characters. Only alphanumeric characters, hyphens (-), and underscores (_) are allowed, however, the name cannot start with a hyphen or an underscore.
7. Select the volume **Emulation**.
8. Select a **Tier** to add to the policy and then specify a storage group capacity for the tier (% MAX of Storage Group). This value is the maximum amount (%) of the associated storage group's logical capacity that the FAST controller can allocate to the tier. This value must be from 1 to 100. The total capacities for a policy must equal to or be greater than 100.
9. Repeat the previous step for any additional tiers you want to add.
10. Click **OK**.

Renaming FAST policies

This procedure explains how to rename FAST policies.

Before you begin:

Policy names must be unique and cannot exceed 32 characters. Only alphanumeric characters, hyphens (-), and underscores (_) are allowed, however, the name cannot start with a hyphen or an underscore.

To rename a FAST Policy:

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** details view.
5. Select the tier and click **View Details** to open the policy's **Details** view.
6. Type a new name for the policy.
7. Click **Apply**.

Deleting FAST policies

This procedure explains how to delete FAST policies.

Before you begin:

You cannot delete a policy that has one or more storage groups associated with it. To delete such a policy, you must first [disassociate the policy from the storage groups](#).

To delete a FAST Policy:

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** list view.
5. Select the policy and click **Delete**.
6. Click **OK** in the confirmation message.

Associating/Disassociating FAST policies and storage groups

The procedure for associating FAST policies and storage groups, depends on whether you are associating a storage group with a policy or policy with a storage group.

Note that you can only disassociate a storage group from a policy, not a policy from a storage group.

Before you begin:

- ◆ Storage groups and FAST policies can only be associated under the following conditions:
 - The storage group is not already associated with another policy. This restriction does not apply to Symmetrix systems running Engenuity 5876.
 - The target FAST policy needs to have a least one pool that is part of the source policy in re-association activity.
 - The volumes in the new storage group are not already in a storage group associated with a FAST policy.
 - The policy has at least one tier.
 - The storage group only contains meta heads; meta members are not allowed.
 - The storage group does not contain moveable volumes. When a storage group is associated with a policy, you cannot add non-moveable volumes to it. Non-moveable volumes include:
 - CKD EAV
 - DRV
 - SFS
 - iSeries, ICOS, ICL
 - SAVE volumes
 - VDEVs
 - Diskless volumes
- ◆ The storage group cannot contain a volume that is part of another storage group already associated with another policy.
- ◆ The Symmetrix system has less than the maximum number of allowed associations (8,192).

To associate a FAST policy with a storage group:

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** details view.
5. Select the policy and click **Associate Storage Group** to open the **Associate Storage Group** dialog box.
6. Select one or more storage groups from the **Available Storage Groups** list and click **Add** to move them to the **Associated Storage Groups** list, or click **Add All** to move all storage groups from the **Available Storage Groups** list to the **Associated Storage Groups** list.
7. Click either of the following:
 - ◆ **OK** to associate the storage group.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To have FAST factor the R1 volume statistics into move decisions made for the R2 volume, select **Enable FAST [VP/DP] RDF Coordination**. This attribute can be set on a storage group, even when there are no SRDF volumes in the storage group. This feature is only available if the Symmetrix system is part of an SRDF setup. Both R1 volumes and R2 volumes need to be running Enginuity version 5876 or higher for the FAST VP system to coordinate the moves. However, the setting of the RDF coordination attribute will not be prevented if one of the Symmetrix systems is running an Enginuity level lower than 5876.
2. Click **OK**.

To associate a storage group with a FAST policy:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Associate to FAST** to open the **Associate to FAST Policy** dialog box.
4. Select a policy and click **OK**.

To disassociate a storage group from a FAST Policy:

1. Select the Symmetrix system.
2. In the navigation bar, click **Storage** to open the **Storage** section.
3. Click **FAST** to open the **FAST** dashboard.
4. Select the **FAST Type**.
5. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** details view.
6. Select the policy and click **View Details** to open the policy's details view.
7. In the **Related Object** view panel, click **Storage Groups** to open the **Storage Groups for FAST Policy** details view.
8. Select the one or more storage groups and click **Disassociate**.
9. Click **OK** in the confirmation message.

Reassociating FAST polices and storage groups

This procedure explains how to reassociate a storage group with a new policy. When reassociating a storage group, all the current attributes set on the original association automatically propagate to the new association. This feature eliminates the previous process of disassociating a storage group, then associating the group to a new policy, and entering the attributes, such as priority, on the association.

Before you begin:

- ◆ The storage group name must be valid.
- ◆ The storage group and policy must already exist on the Symmetrix system.
- ◆ The storage group must be in an association before performing a reassociation.
- ◆ The new policy for the storage group, must have the same emulation as the storage group. Mix emulation association will result in an error.
- ◆ The storage group cannot be associated with an empty policy, and the reassociated policy must contain at least one tier.
- ◆ The total of the capacity percentage for the target FAST policy must add up to at least 100%.
- ◆ If the FAST policy contains VP Tiers, all of the thin devices in the storage group must be bound to any VP pool in a tier in the policy. None of the thin devices can be bound to a pool outside of the policy.

To reassociate FAST Policies and storage groups:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Reassociate to FAST Policy** to open the **Reassociate to FAST Policy** dialog box.
4. Select a policy and click **OK**.

Viewing FAST policies

1. Select the Symmetrix system.
 2. Select **Storage** > **FAST** to open the **FAST** dashboard.
 3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
 4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** list view.
- Use the **FAST Policies** list view to view and manage FAST policies on a Symmetrix system.

The following properties display:

- ◆ **Policy Name** — Name of the policy.
- ◆ **Tier 1** — Symmetrix tier associated with the policy.
- ◆ **Tier 1 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 1.
- ◆ **Tier 2** — Symmetrix tier associated with the policy.
- ◆ **Tier 2 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 2.
- ◆ **Tier 3** — Symmetrix tier associated with the policy.
- ◆ **Tier 3 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 3.
- ◆ **Tier 4** — Symmetrix tier associated with the policy.
- ◆ Up to 4 tiers is supported only for FAST VP policies. FAST policies support up to 3 tiers.
- ◆ **Tier 4 %** — Maximum amount (%) of an associated storage group that can be allocated to the Symmetrix tier 4.
- ◆ **# Associated Storage Group** — Storage group associated with the policy.

The following controls are available:

- ◆ **Create** — See [FAST policies](#) on page 180.
- ◆ **View Details** — See [Viewing FAST policy details](#) on next page.
- ◆ **Delete** — See [Deleting FAST policies](#) on page 181.
- ◆ **Associated Storage Groups** — See [Associating/Disassociating FAST policies and storage groups](#) on page 183.

Viewing FAST policy details

1. Select the Symmetrix system.
2. Select **Storage** > **FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Policies** view block, click **Manage Policies** to open the **FAST Policies** list view.
5. Select the policy and click **View Details** to open its **Details** view.

The policy **Details** view allows you to view and manage a FAST Policy. It contains [Properties](#), [Related Objects](#), [Performance Views](#), and [Graphs](#) panels.

Properties panel

The following properties display:

- ◆ **Policy Name** — Name of the policy. To rename the policy, type a new name over the existing and click **Apply**. Policy names must be unique and cannot exceed 32 characters. Only alphanumeric characters, hyphens (-), and underscores (_) are allowed, however, the name cannot start with a hyphen or an underscore.
- ◆ **Tier 1 - 3** (for FAST DP)
- ◆ **Tier 1 - 4** (for FAST VP) — Symmetrix tier associated with the policy, followed by the maximum amount (%) of the associated storage group's logical capacity that the FAST controller can allocate to the tier. This value must be from 1 to 100. The total capacities for a policy must be greater than or equal to 100. To change a tier, select another from the list, and click **Apply**. To change a maximum amount, type a new amount, and click **Apply**.

The following controls are available:

- ◆ **Associate Storage Groups** — See [Associating/Disassociating FAST policies and storage groups](#) on page 183.
- ◆ **Expand** — See [Expanding storage groups](#) on page 152.
- ◆ **Delete** — See [Deleting storage groups](#) on page 162.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the policy.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Objects panel

The **Related Objects** panel links you to views displaying objects contained in and associated with the FAST Policy. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Tiers - 3** will open a view listing the three tiers in the policy.

Performance Views panel

The **Performance** panel links you to the performance monitor and analyze views for the policy.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Graphs panel

The **Graphs** panel includes graphic representations of the used and free space available for each tier in the policy. In addition, each chart includes markers for the following metrics:

- ◆ **Max SG Demand** — The calculated upper limit for the storage group on the tier.
- ◆ **Available to FAST** — The amount of storage available for FAST operations on the tier.

FASTcontroller

Setting FAST control parameters

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Status Report** view block, click **Edit** (next to **Settings**) to open the corresponding **FAST Settings** dialog box.
5. Modify any number of the following parameters and click **OK**. Note that the parameters available to you depend on the version of FAST:

Table 5: Comparison of FAST DP and Fast VP

FAST DP	FAST VP
<p>Set State — Enables/Disables the FAST controller. When Enabled, the FAST controller will move data between tiers based on the defined policies. When Disabled, the FAST controller will not move data between tiers. By default, the FAST controller is disabled.</p>	
<p>Data Movement Mode — Sets the mode of the FAST controller to automatic or user approval mode. If the FAST controller is set to user approval mode, it will generate plans, but not perform any movements unless the plans are approved by the user. This option is shared with Symmetrix Optimizer.</p>	
<p>Maximum Moves Per Day — Specifies the maximum number of moves to perform in a 24 hour period, starting at 12:00 AM. Possible values range from 2 to 200, with 200 being the default. This option is shared with Symmetrix Optimizer.</p>	<p>Relocation Rate — Specifies the data movement mode for thin volumes. If set to Automatic, the FAST system will continuously perform data movement for thin volumes within the data movement window, without user intervention. If set to Off, the FAST controller will not perform any data movement for thin volumes. There is no equivalent to user approval mode for thin data movement.</p>
<p>Maximum Simultaneous Moves — Specifies the maximum number of moves that can be performed at one time. Possible values range from 2 to 32, with 32 being the default. This</p>	<p>Pool Reserved Capacity — Specifies the capacity of each thin pool (percentage) that will be reserved for non-FAST activities. If the free space in a given thin pool (as a percentage of pool-enabled capacity) falls below this value, the FAST controller will not move any more chunks (a group of 12 tracks) into that pool. To move any new chunks to the pool, the FAST controller must first move some chunks from that pool to another pool to free</p>

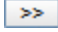
FAST DP	FAST VP
<p>option is shared with Symmetrix Optimizer.</p>	<p>up space. Enforcement of this parameter is best-effort; FAST may move chunks to a thin pool resulting in a violation because non-FAST activities (such as new allocations for writes to a thin volume) can simultaneously consume pool free capacity. Possible values range from 1 to 80, with 10 being the default.</p>
<p>Migration Restriction — Specifies whether the FAST controller can perform swaps and moves, or only swaps.</p>	<p>Allocate by FAST Policy — When enabled, the system chooses a pool for the policy when making an allocation for a thin volume.</p>
<p>Allow FAST to use volumes that are not visible to the host for full swaps — Indicates if the FAST controller can use host invisible volumes (unmasked and unmapped) to do a full swap with volumes in storage groups under FAST control to improve the performance of the storage group. Possible values are ENABLE and DISABLE. The default is DISABLE.</p>	<p>Compression Attributes— Enables FAST VP Compression. When enabled, the FAST controller will compress the members of any thin pool (for which this feature is enabled) according to the following attributes:</p> <ul style="list-style-type: none"> ◆ Time to Compress — Specifies how frequently compression will occur. Valid values are 1 - 365 days. Default is 40 days. ◆ Compression Rate — Specifies the compression rate. Valid values are 1 (most aggressive) to 10 (least aggressive). Default is 5. <p>You can enable this feature for a thin pool while creating it (Thin pools (page 231) or post creation in the pool's Details view (Viewing thin pool details (page 240)).</p> <p>In addition, you can also manually control this feature at the storage group level (Managing VP compression on storage groups (page 163)).</p>

Pinning and unpinning volumes

Pinning volumes prevents any automated process such as FAST or Optimizer from moving them. However, you can still manually move a pinned volume with Optimizer or migrate a pinned volume with Virtual LUN Migration.

Note The capacity of pinned volumes is counted for compliance purposes.

To pin and unpin volumes:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Select one or more volumes, click more , and select one of the following:
 - ◆ **Pin** — To pin the volumes.
 - ◆ **Unpin** — To unpin the volumes.
4. Click **OK** in the confirmation message.

Time windows

Understanding time windows

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

Note EMC® Unisphere for VMAX supports time windows in the enhanced format introduced in Enginuity 5875 Q2 2011 SR. To use time windows created in earlier versions of Enginuity, you must convert them according to the procedure in the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

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.

Creating and modifying time windows

This procedure explains how to create/modify time windows.


Before you begin:

- ◆ Time windows are used by FAST and Optimizer. Changes made to FAST time windows may also affect Optimizer.
- ◆ The maximum number of time windows that can be defined on a Symmetrix system is 128.

To create time windows:

1. Select the Symmetrix system.
2. Select **Storage > FAST** to open the **FAST** dashboard.
3. If the Symmetrix system is licensed for both FAST DP and FAST VP, select the **FAST Type** to which the time window will apply.
4. In the **FAST Status Report** view block, click **Edit** next to the type of time window you want to create/modify.

Depending on your selection, either the **FAST Performance Time Window** or the **FAST Movement Time Window** dialog box opens.

5. In the **Show** field, click the calendar icon  and select the week in which to define the time window.
6. Click **Show Advanced**.
7. Do the following, depending on whether you are creating/modifying an open or closed time window:

Open time window:

- a. Click **Manage** next to the **Open Time Window (Inclusive)** option to open **Manage Open Time Window** dialog box.
- b. Define the time window, by selecting one of the following options and clicking **Add**:
 - ◆ **Always open** — Creates a single open time window for the entire week (Sunday to Saturday).
 - ◆ **All weekend** — Creates a single open time window for the weekend (17:00 Friday to 8:00 Monday).
 - ◆ **9:00 to 17:00 , Monday to Friday** — Creates five time windows, one for each day of the work week.
 - ◆ **17:00 to 8:00, Monday to Friday** — Creates five time windows, one for each of night of the work week.
 - ◆ **Custom** — Allows you to define your own time window.
- c. Click **OK** to close the **Manage** dialog box.

Closed time window:

- a. Click **Manage** next to the **Closed Time Window (Exclusive)** option to open **Manage Closed Performance Time Window** dialog box.
- b. Define the time window, by selecting a **Start Date/Time** and an **End Date/Time**, and clicking **Add**.
- c. Click **OK** to close the **Manage** dialog box.

8. Define the following parameters:
 - ◆ **Workload Analysis Period** — Specifies the amount of workload sampling to maintain for sample analysis. Possible values are specified in units of time (hours, days, or weeks) and can range from 2 hours to 4 weeks, with the default being 8 hours.
 - ◆ **Time to Sample Before First Analysis** — Specifies the minimum amount of workload sampling to complete before analyzing the samples for the first time. When setting this parameter, be sure to allow enough time (usually a week) to establish a good characterization of the typical workload. This parameter allows you to begin operations before the entire Workload period has elapsed. Possible values range from 1 hour to the value specified for the Workload Analysis Period parameter.
9. Click **OK**.

Volume configuration

Creating regular volumes

This procedure explains how to create regular volumes from the **Volume Dashboard**. You can also create them by clicking the **Create Volumes** option in the **Common Tasks** panel.

To create regular volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box. The dialog box opens with **Regular** volume option selected by default.
5. Select whether to create **Standard** volumes or **BCVs**.
6. Select the **Disk Technology**.
External disk technology is an option if the Symmetrix system has FTS (Federated Tiered Storage) enabled and available external storage.
7. Select the **Emulation** type.
8. Select the RAID **Protection** level.
9. Specify the capacity to create by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
10. Click one of the following:
 - ◆ **Add to Job List** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To create the volumes from a specific **Disk Group**, select one (disk group number and name).
2. To assign **Dynamic Capability** to the volumes, select one of the following; otherwise, leave this field set to **None**.
 - ◆ **RDF1_Capable** — Creates a dynamic R1 RDF volume.
 - ◆ **RDF2_Capable** — Creates a dynamic R2 RDF volume.
 - ◆ **RDF1_OR_RDF2_Capable** — Creates a dynamic R1 or R2 RDF volume.
3. View **Enable SCSI3 Persistent Reservation** status — For Enginuity 5875 and higher this feature is pre-set by SYMAPI and cannot be changed. It is displayed as an it is enabled for Enginuity 5875 and higher. It is displayed as enabled for Enginuity 5875 and higher, except for CDK and AS/400 emulations.
4. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name**:
 - ◆ **Name Only** — All volumes will have the same name.

- ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
- ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

For more information on naming volumes, refer to [Setting volume names](#) on page 212

5. In the **Define Meta** panel, enter meta configuration parameters:
 - ◆ **Member capacity (Cyl/MB/GB)**
 - ◆ **Configuration (Striped/Concatenated)**

If auto meta is enabled on the system then it displays as enabled with a green check mark.

The **Define Meta** panel only displays when attempting to create a volume larger than the value specified in the Minimum Auto Meta Size.

6. Click either of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating DATA volumes

This procedure explains how to create DATA volumes from the **Volume Dashboard**. You can also create them by clicking the **Create volumes** option in the **Common Tasks** panel.

Before you begin:

DATA volumes are used with Symmetrix Virtual Provisioning, which requires Engenuity 5773 or higher.

To create DATA volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Private**.
6. Select **DATA** as the **Configuration**.
7. Select the **Disk Technology**.

External disk technology is an option if the Symmetrix system has FTS (Federated Tiered Storage) enabled and available external storage.

8. Select the **Emulation** type.
9. Select the RAID **Protection** level.

10. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
11. To add the new volumes to a specific thin pool, select one from **Add to Pool**. Pools listed are filtered on technology, emulation, and protection type.
12. Click one of the following:
 - ◆ **Add to Job List** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. Select the **Disk Group** (number and name) in which to create the volumes. The list of disk groups is already filtered based on technology type selected above.
2. To enable the new volumes in the pool, select **Enable volume in pool**.
3. Click either of the following:
 - ◆ **Add to Job List** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating SAVE volumes

This procedure explains how to create SAVE volumes from the **Volume Dashboard**. You can also create them by clicking the **Create Volumes** option in the **Common Tasks** panel.

SAVE volumes are used to store pre-update data from TimeFinder/Snap source devices or to store excess cycle data from SRDF/A Delta Set Extension (DSE) operations.

To create SAVE volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Private**.
6. From **Configuration**, select **SAVE**. When creating SAVE volumes on Symmetrix systems running Enginuity 5875 for VMAXe, this field is set to **2-Way Mir**, by default.
7. Select the **Disk Technology**.

External disk technology is an option if the Symmetrix system has FTS (Federated Tiered Storage) enabled and available external storage.
8. Select the **Emulation** type.
9. Select the RAID **Protection** level.
10. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.

11. To add the new volumes to a specific pool, select one from **Add to pool**. SNAP and SRDF/A DSE pools listed are filtered on technology, emulation, and protection type selected above.
12. Click one of the following:
 - ◆ **Add to Job List**. to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

If Auto meta is enabled on the system then it displays as enabled with a green check mark.

Setting Advanced options:

1. Select the **Disk Group** (number and name) in which to create the volumes. The list of disk groups is already filtered based on technology type selected above.
2. To enable the new volumes in the pool, select **Enable volume in pool**.
If Auto meta is enabled on the system then it displays as enabled with a green check mark.
3. Click one of the following:
 - ◆ **Add to Job List**. to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating thin volumes

This procedure explains how to create thin volumes from the **Volume Dashboard**. You can also create them by clicking the **Create Volumes** option in the **Common Tasks** panel.

To create thin volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Virtual**.
6. Select **Configuration (TDEV or BCV + TDEV)** or thin volumes.
7. Select the **Emulation** type.
8. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
9. To bind the new volumes to a specific thin pool, select one from **Bind to Pool**. Only thin pools with enabled DATA volumes and matching emulation are available for binding (except AS/400 which will bind to an FBA pool).

10. Click one of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. If creating a thin volume or a thin BCV, you can specify the amount of space to preallocate on each volume, by selecting **Preallocate capacity for each new volume** selecting either **Full Volume** or **By Capacity**, and then typing a value. The amount of space must be less than or equal to the available space in the pool and less than or equal to the allocated size of the thin volume. This option is not available when creating VDEVs.
2. If you specified to preallocate space in the previous step, you can mark the preallocation on the thin volume as persistent by selecting **Persist preallocated capacity through reclaim or copy**. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations. This option is not available when creating VDEVs.
3. To assign **Dynamic Capability** to the volumes, select one of the following; otherwise, leave this field set to **None**.
 - ◆ **RDF1_Capable** — Creates a dynamic R1 RDF volume.
 - ◆ **RDF2_Capable** — Creates a dynamic R2 RDF volume.
 - ◆ **RDF1_OR_RDF2_Capable** — Creates a dynamic R1 or R2 RDF volume.
4. View **Enable SCSI3 Persistent Reservation** status — For Enginuity 5875 and higher this feature is pre-set by SYMAPI and cannot be changed. It is displayed as enabled for Enginuity 5875 and higher, except for CDK and AS/400 emulations.
If Auto meta is enabled on the system then it displays as enabled with a green check mark.
5. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name**:
 - ◆ **Name Only** — All volumes will have the same name.
 - ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

For more information on naming volumes, refer to [Setting volume names](#) on page 212

6. Click one of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating gatekeeper volumes

This procedure explains how to create gatekeeper volumes from the **Volume Dashboard**. You can also create them by clicking the **Create Volumes** option in the **Common Tasks** panel.

To create gatekeeper volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Private**.
6. Select **Gatekeeper** or **Virtual Gatekeeper** as the **Configuration**.
7. Select the **Emulation** type.
8. Type the **Number of Volumes** to create.
9. If creating virtual gatekeepers, you can optionally bind them to a specific thin pool (**Bind to Pool**). Only thin pools with enabled DATA volumes and matching emulation are available for binding (except AS/400 which will bind to an FBA pool).
10. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating diskless volumes

This procedure explains how to create diskless volumes from the **Volume Dashboard**. You can also create them by clicking the **Create volumes** option in the **Common Tasks** panel.

To create DLDEV volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Private**.
6. Select **DLDEV** as the **Configuration**.
7. Select the **Emulation** type.
8. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter the volume capacity.

9. Click one of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To assign **Dynamic Capability** to the volumes, select one of the following; otherwise, leave this field set to **None**.
 - ◆ **RDF1_Capable** — Creates a dynamic R1 RDF volume.
 - ◆ **RDF2_Capable** — Creates a dynamic R2 RDF volume.
 - ◆ **RDF1_OR_RDF2_Capable** — Creates a dynamic R1 or R2 RDF volume.

The **Define Meta** panel only displays when attempting to create a volume larger than the value specified in the Minimum Auto Meta Size.

2. In the **Define Meta** panel, enter meta configuration parameters:

- ◆ **Member capacity (Cyl/MB/GB)**
- ◆ **Configuration (Striped/Concatenated)**

If auto meta is enabled on the system then it displays as enabled with a green check mark.

3. Click either of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating DRV volumes

This procedure explains how to create DRV volumes from the **Volume Dashboard**. You can also create them by clicking the **Create volumes** option in the **Common Tasks** panel.

To create DRV volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Private**.
6. From **Configuration**, select **DRV**.
7. Select the **Disk Technology** on which to create the volumes.
8. Select the **Emulation** type.
9. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.

10. Click one of the following:

- ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
- ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
- ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To create the volumes from a specific disk group, select one (disk group number and name) from **Disk Group**.

If Auto meta is enabled on the system then it displays as enabled with a green check mark.

2. Click one of the following:

- ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
- ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating volumes using storage templates

This procedure explains how to create volumes from the **Volume Dashboard**. You can also create them by clicking the **Create Volumes** option in the **Common Tasks** panel.

To create volumes using storage templates:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Template**.
6. Type the name of the **Template** to use when creating the volumes, or click **Select** to open a dialog box from which you can select a template.
7. Type the **Number of Volumes** to create. This field appears populated or dimmed if the template you are using contains a total capacity value.

If Auto meta is enabled on the system then it displays as enabled with a green check mark.

8. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name**:
 - ◆ **Name Only** — All volumes will have the same name.
 - ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

For more information on naming volumes, refer to [Setting volume names](#) on page 212

2. Click either of the following:
 - ◆ **Add to Job List**. to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Deleting volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Select the volume and click **Delete**.
5. Click **Delete** in the confirmation dialog box.


Duplicating volumes

The following explains how to duplicate volumes.

Before you begin:

- ◆ When attempting to duplicate thin volumes, only thin volumes configured as a TDEV can be duplicated. If you are duplicating a thin volume that is bound to a pool, the newly created volumes will be bound to the same pool.
- ◆ If you are duplicating a DATA volume that is part of a pool, the newly created DATA volumes will be part of the same pool. The initial state of the volume will be DISABLED.

To duplicate a volume:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, select the type of volume and click **View** to open the **Volumes** list view.
4. Select the volume, click more , and click **Duplicate** to open the **Duplicate Volume** dialog box.
5. Type the **Number of Volumes** (duplicates) to make.
6. *z/OS Only:* You can optionally change the SSID number for the new volumes by typing a new value, or clicking **Select...** to open a dialog from which you can select an SSID. By default, this field displays the SSID of the volume you are copying.
7. Click one of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name**:
 - ◆ **Name Only** — All volumes will have the same name.
 - ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

For more information on naming volumes, refer to [Setting volume names](#) on page 212

2. Do one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and


[Previewing/Running jobs](#) on page 99.

- ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

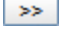
Assigning Symmetrix priority

You can assign Symmetrix priority to individual volumes or groups of volumes (DGs or SGs).

To assign Symmetrix priority to individual volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, select the type of volume.
4. Click **View** to open its list view.
5. Select one or more volumes, click more , and select **Assign Symmetrix Priority** to open the **Assign Symmetrix Priority** dialog box.
6. *Optional:* Click **Show Selected Volumes** to view details on the selected volumes.
7. Select a **Symmetrix Priority** from 1 (the fastest) to 16 (the slowest) and click **OK**.
8. Click **OK** in the confirmation message.

To assign Symmetrix priority to groups of volumes:

1. Select the Symmetrix system.
2. Select the group (**Storage > Storage Groups**, or **Data Protection > Device Groups**).
3. Click more  and select **Assign Symmetrix Priority** to open the **Assign Symmetrix Priority** dialog box.
4. *Optional:* Click **Show Selected Volumes** to view details on the selected volumes.
5. Select a **Symmetrix Priority** from 1 (the fastest) to 16 (the slowest) and click **OK**.
6. Click **OK** in the confirmation message.

Changing volume configuration


This procedure explains how to change a volume's configuration.

Before you begin:

- ◆ On Symmetrix systems running Enginuity 5874, you cannot increase or decrease the mirror protection of a volume. For example, you cannot convert a 2-way mirror to a 4-way mirror, and you cannot convert a 4-way mirror to a 2-way mirror. However, on Symmetrix systems running Enginuity 5773, you can increase or decrease the mirror protection of a volume.
- ◆ On Symmetrix systems running Enginuity 5874, you can convert TDEVs to BCVs and TDEV to BCVs to TDEVs. No other conversions to/from TDEV BCVS are allowed.

To change the configuration of volumes:

1. Select the Symmetrix system.

2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, double-click the type of volume on which to set the attributes. The **Volumes** list view opens.
4. Select one or more volumes, click more  , and click **Change Volume Configuration** to open the **Change Volume Configuration** dialog box.
5. To view details on the selected volumes, click **Show Selected Volumes**.
6. Select a **New Configuration** for the selected volumes. Only valid configurations are listed. The remaining fields in the dialog box are active or inactive depending on the configuration type.
7. *z/OS Only:* Type the **SSID** for the new volume created by removing a mirror, or click **Select...** to open a dialog from which you can select an SSID. This is required for volumes on Symmetrix systems with ESCON or FICON directors (or mixed systems). This field is optional on Symmetrix systems running Engenuity 5773 or higher when reducing the number of mirrors.
8. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
9. Click **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.

Enabling and disabling volumes

This procedure explains how to enable or disable volumes for use in a pool (Snap, SRDF/A DSE, Thin). The volumes in the pool do not all have to be in the same state (enabled or disabled). If all the volumes in a pool are disabled, then the pool is disabled. If at least one volume in a pool is enabled, then the pool is enabled.

Before you begin:


To disable a volume, all sessions must be terminated, and have no used tracks.

To enable/disable DATA volumes:


1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select either **SAVE** or **DATA**. The corresponding list view opens.
4. To enable volumes, right-click them and select **Enable**. To disable volumes, right-click them and select **Disable**.
5. Click **OK** in the confirmation message.

Mapping volumes

1. Select the Symmetrix system.

2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, select the type of volume.
4. Click **View** to open its list view.
5. Select one or more volumes, click more , and select **Map** to open the mapping wizard.
6. Select one or more **Available** volumes and click **Add** to move them to the **Selected** volumes list, or click **Add All** to move all **Available** volumes to the **Selected** volumes list.
7. Click **Next**.
8. Select one or more **Available** ports and click **Add** to move them to the **Selected** ports list, or click **Add All** to move all **Available** ports to the **Selected** ports list.
9. Click **Next**.
10. To change an automatically generated LUN address, do the following; otherwise, click **Next** to accept the generated address.
 - a. Double-click an address to open the **Set LUN Address** dialog box.
 - b. Double-click an address to type a new address over it, or select an address and click **Next Available** to increment the generated address to the next available address. To reset an address back to the generated address, click **Reset**.
 - c. Click **OK** to return to the mapping wizard
 - d. Click **Next**.
11. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
12. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Unmapping volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, select the type of volume.
4. Click **View** to open its list view.
5. Select one or more volumes, click more , and select **Unmap** to open the unmap wizard.
6. Select one or more **Available** volumes and click **Add** to move them to the **Selected** volumes list, or click **Add All** to move all **Available** volumes to the **Selected** volumes list.
7. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in [Symmetrix Preferences](#) for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
8. Click **Next**.
9. Select one or more **Available** ports and click **Add** to move them to the **Selected** ports list, or click **Add All** to move all **Available** ports to the **Selected** ports list.
10. Click **Next**.
11. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
12. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Setting optimized read miss

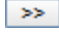
The optimized read miss feature reduces I/O processing overhead of read miss operations for both DA and DX emulations. The feature is supported on Symmetrix VMAX arrays running Engenuity 5876 Q2 2013 or higher.

This procedure explains how to set the optimized miss feature at the volume level. You can also perform this operation at the storage group or the device group level

Before you begin:

The optimized read miss feature is supported only for EFD volumes with FBA or AS400 D910 emulation attached to an XtremSW Cache Adapter.

To set cacheless read miss:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** dashboard.
3. In the **Volume Type** panel, select the type of volume.
4. Click **View** to open its list view.
5. Select one or more volumes, click more  , and select **Set Optimized Read Miss** to open the **Optimized Read Miss** dialog box.
6. *Optional:* Click **Show Selected Volumes** to view details on the selected volumes.
7. Select an **Optimized Read Miss** mode:
 - ◆ **Off** — Disables cacheless read miss mode
 - ◆ **System Managed** — Symmetrix system determines the appropriate cacheless read miss mode for the specified volumes.
8. Click **OK**.


Setting volume status

This procedure explains how to change the status of volumes.

Before you begin:

You cannot set the status of an unbound thin volume.

To set volume status for individual volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In **Volume Type**, select the type of volume on which to set the identifier.
4. Click **View** to open the volume list view.
5. Select one or more volumes, click more , and select **Set Volume Status** to open the **Set Volume Status** dialog box.
6. To view details on the selected volumes, click **Show Selected Volumes**.
7. Select a **Status**. Possible values are:
 - ◆ **Read/Write Enable** — Changes the write-protect status of the volumes to be read and write enabled on the specified director port(s) for any locally attached hosts.
 - ◆ **Write Disable** — Changes the write-protect status of the volumes to be write disabled on the specified director ports for any locally attached hosts. This option will only work on volumes that are in a write enabled state.
 - ◆ **Read/Write Enable - on all Ports** — Changes the write-protect status of the volumes to be read and write enabled on all the director ports for any locally attached hosts.
 - ◆ **Write Disable - on all Ports** — Changes the write-protect status of the volumes to be write disabled on all the director ports for any locally attached hosts.
 - ◆ **Device Ready** — Changes the User Ready status of the volumes to Ready.
 - ◆ **Device Not Ready** — Changes the User Ready status of the volumes to Not Ready.
 - ◆ **Hold** — Causes the Hold bit to be placed on a volume. The Hold bit is automatically placed on the target volume of a Snap session.
 - ◆ **Unhold** — Causes the Hold bit to be removed from a volume. The Hold bit is automatically removed from the target volume of a snap session when the snap session is removed.
8. If the selected volumes are mapped, you can select to change the status for a particular director or all directors.
9. Click **OK**.


Setting volume attributes

Setting the volume attribute for a volume restricts how it can be accessed.


Before you begin:

You cannot set attributes for DATA volumes.

To set volume attributes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In **Volume Type**, select the type of volume on which to set the attributes.
4. Click **View** to open the volume list view.
5. Select one or more volumes, click more , and select **Set Volume Attributes** to open the **Set Volume Attributes** dialog box.
6. To view details on the selected volumes, click **Show Selected Volumes**.
7. Set any number of the following attributes. Note that the attributes available depend on the type of selected volumes.
 - ◆ **Emulation** — Sets the emulation type for the volumes. The default is No Change. This option will appear dimmed for masked/mapped volumes, as this feature is not supported on masked/mapped volumes.
 - ◆ **Dynamic RDF Capability** — Sets the volume to perform dynamic RDF operations. Possible operations are:
 - **No Change** — Keeps the RDF capability the same.
 - **None** — Sets the volume for no dynamic RDF capability.
 - **RDF1 or RDF2 Capable** — Allows the volume to be R1 or R2 (RDF swaps allowed). Select this attribute to create an R21 volume used in a Cascaded RDF operation.
 - **RDF1 Capable** — Allows the volume to be an R1 (no RDF swaps).
 - **RDF 2 Capable** — Allows the volume to be an R2 (no RDF swaps).
 - ◆ **SCSI3 Persistent Reservation** — Maintains any reservations (flags) whether the system goes online or offline. This field will appear dimmed for diskless volumes.
8. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
9. Click **OK**.

Setting volume identifiers

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In **Volume Type**, select the type of volume on which to set the identifier.
4. Click **View** to open the volume list view.
5. Select one or more volumes, click more , and select **Set Volume Identifiers** to open the **Set Volume Identifiers** dialog box.
6. To view details on the selected volumes, click **Show Selected Volumes**.
7. Type the **Volume Identifier Name**. Volume identifier names must be unique from other volumes on the Symmetrix system and cannot exceed 64 characters. Only alphanumeric characters and underscores (`_`) are allowed.
8. Type the **Volume HP Identifier Name**. HP identifier names must be a user-defined volume name (not to exceed 128 alpha-numeric characters) applicable to HP-mapped volumes. This value is mutually exclusive of the VMS ID. This attribute will appear grayed out for diskless volumes.
9. Type the **VMS Identifier Name**. VMS identifier names must be a numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID. This attribute will appear grayed out for diskless volumes.
10. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
11. Click **OK**.

Setting volume names

When creating or duplicating volumes; or creating or expanding storage groups, you can optionally name the new volumes.

When naming volumes, you should be aware of the following:

- ◆ Volume names cannot exceed 64 characters. Only alphanumeric characters, underscores (`_`), and periods (`.`) are allowed.
- ◆ Volume names plus an optional suffix cannot exceed 64 characters. If using a numerical suffix, volume names cannot exceed 50 characters (prefix) and the trailing numerical suffix number cannot exceed 14 characters. If not using a numerical suffix, all 64 characters can be specified for the volume name. The maximum starting suffix is 1000000.
- ◆ This feature is not supported for the following types of volumes: SFS, DRV, Meta members, SAVE, DATA, Vault, and diskless.

Managing Meta Volumes

Creating meta volumes

Before you begin:

- ◆ On Symmetrix systems running Enginuity versions lower than 5875, only unmapped thin volumes (hypers) can be formed into meta volumes.
- ◆ On Symmetrix systems running Enginuity 5875 or higher:
 - Bound thin volumes can be used as meta heads; however, bound thin volumes cannot be used as meta members.
 - Unmapped thin volumes can be formed into striped meta volumes.
 - Mapped or unmapped thin volumes can be formed into concatenated meta volumes.
- ◆ For a complete list of restrictions and recommendations on creating meta volumes, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

To create meta volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Meta Volume** folder and double-click a meta type to open the **Meta Volumes** list view.
4. Click **Create** to open the **Create Meta Volume** wizard.
5. Select the **Emulation** type.
6. If creating FBA volumes, select whether to create them from **New** or **Existing** volumes.
7. If creating FBA or AS/400 volumes, select the **Meta Volume Configuration (Concatenated or striped)**.
8. Select a method for forming the meta volumes.
9. Click **Next**.

10. Do the following, depending on the method you selected:

Using existing virtual volumes:

- a. Type the **Number of Meta Volumes** to form.
- b. Specify the **Meta Volume Capacity** by typing the **Meta Volume Member Count including Head**, and selecting a **Meta Volume Member Capacity**.
- c. Select a **Volume Configuration** for the members.
- d. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in [Symmetrix Preferences](#) for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
- e. If you are creating CKD meta volumes, type or select an **SSID**.
- f. If you are creating striped meta volumes, you can optionally select the size of the meta volumes, by clicking **Show Advanced**, and selecting a **Striped Size**. The striped size can be expressed in blocks or cylinders. Possible sizes in 512 byte blocks are 1920, 3840, 7680, 15360, 30720, and 61440. The stripe size must be 1920, which is the default for all versions of Enginuity. If no stripe size is specified when creating a striped meta, all Enginuity codes will consider the default stripe size as 1920 blocks of 512 bytes each.
- g. Click **Next**.

Using existing standard provisioned volumes:

- a. Type the **Number of Meta Volumes** to form.
- b. Specify the **Meta Volume Capacity** by typing the **Meta Volume Member Count including Head**, and selecting a **Meta Volume Member Capacity**.
- c. Select a **Volume Configuration**.
- d. Select the RAID **Protection** level for the meta volumes.
- e. Select the type of **Disk Technology** on which the meta volumes will reside.
- f. Select the **Disk Group** containing the meta volumes.
- g. Click **Next**.

By manually selecting existing volumes:

- a. Select from the listed volumes.
- b. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in [Symmetrix Preferences](#) for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
- c. Click **Next**.

Creating standard provisioned volumes:

- a. Specify the **Meta Volume Capacity** by typing the **Number of Meta Volumes**, and selecting a **Meta Volume Capacity**.
- b. Select a **Volume Configuration**.
- c. Select the RAID **Protection** level for the meta volumes.
- d. Select the type of **Disk Technology** on which the meta volumes will reside.

- e. If you are creating CKD meta volumes, type or select an **SSID**.
11. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
 12. Click **Run Now**.

Adding meta members

Before you begin:

To expand a bound striped thin meta volume on a Symmetrix system running an Enginuity version lower than 5875, you must first [unbind the volume from the pool](#). This operation is allowed on Symmetrix systems running Enginuity 5875 or higher without having to unbind the volume, however, you must select the **Protect Data** option.

To add meta members:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the type of meta volume you want to expand. The **Meta Volumes** list view opens.
4. Select the meta volume and click **Add Member** to open the **Add Member** dialog box.
5. *For striped metas only:* To protect the original striped meta data, do the following:
 - a. Select the **Protect Data** option.
 - b. Type or select the name of the BCV meta head to use when protecting the data. By default, this field is filled in with the first available BCV.
6. Select one or more volumes to add to the meta volume.
7. Click **Add to Job List** and refer to [Managing job lists](#) on page 97.

Removing meta members

Before you begin:

You can only remove members from concatenated meta volumes.

To remove meta members

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the meta volume type to open the **Meta Volumes** list view.
4. Select the meta volume and click **View Details** to open its **Details** view.
5. In the **Related Objects** panel, click **Meta Members** to open the **Meta Members** list view.
6. Select one or more members and click **Remove Member** to open the **Remove Meta Volume Member** dialog box.
7. Click **Run Now**.

Converting meta volumes

This procedure explains how to change the configuration of a meta volume:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the type of meta volume you want to convert. The **Meta Volumes** list view opens.
4. Select the meta volume and click **Convert** to open the **Convert Meta Volume** dialog box.
5. If converting from concatenated to striped, you can optionally specify to protect the original striped data by selecting **Protect Data** and typing or selecting the BCV meta head to use when protecting the data. By default, the **BCV** field is filled in with the first available BCV.
6. Click **Add to Job List** and refer to [Managing job lists](#) on page 97.

Dissolving meta volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the type of meta volume you want to dissolve. The **Meta Volumes** list view opens.
4. Select the meta volume and click **Dissolve** to open the **Dissolve Meta Volume** dialog box.
5. Click **Add to Job List** and refer to [Managing job lists](#) on page 97.

Viewing meta volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the meta type to open the **Meta Volumes** list view.

Use the **Meta Volumes** list view to view and manage meta volumes.

The following properties display:

- ◆ **Name** — Symmetrix volume number.
- ◆ **Type** — Volume type.
- ◆ **Meta Config** — Volume configuration.
- ◆ **Striped Size** — Meta striped size.
- ◆ **Status** — Volume status.
- ◆ **Reserved** — Whether the volume is reserved.
- ◆ **Capacity (GB)** — Volume capacity in GB.
- ◆ **Emulation** — Emulation type.
- ◆ **Paths** — Number of masking records for the volume.

The following controls are available:

- ◆ **Create** — See [Managing Meta Volumes](#) on page 213.
- ◆ **View Details** — See [Viewing meta volume details](#) on next page.
- ◆ **Add Member** — See [Adding meta members](#) on page 215.
- ◆ **Dissolve** — See [Dissolving meta volumes](#) on previous page.
- ◆ **Convert** — See [Converting meta volumes](#) on previous page.
- ◆ **VLUN Migration** — See [Migrating regular volumes](#) on page 385.
- ◆ **Unpin** — See [Pinning and unpinning volumes](#) on page 191.
- ◆ **Pin** — See [Pinning and unpinning volumes](#) on page 191.
- ◆ **Assign Symmetrix Priority** — See [Assigning Symmetrix priority](#) on page 205.
- ◆ **Assign Dynamic Cache Partition** — See [Assigning volumes to dynamic cache partitions](#) on page 131.
- ◆ **Set Volume Status** — See [Setting volume status](#) on page 210.
- ◆ **Set Volume Attributes** — See [Setting volume attributes](#) on page 211.
- ◆ **Change Volume Configuration** — See [Changing volume configuration](#) on page 205.
- ◆ **Unmap** — See [Unmapping volumes](#) on page 208.

- ◆ **Map** — See [Mapping volumes](#) on page 206.
- ◆ **Untag for RecoverPoint** — See [Tagging and untagging volumes for RecoverPoint](#) on page 410.
- ◆ **Tag for RecoverPoint** — See [Tagging and untagging volumes for RecoverPoint](#) on page 410.

Viewing meta volume details

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the meta type to open the **Meta Volumes** list view.
4. Select a volume and click **View Details** to open its **Details** view.

Use the meta volume **Details** view to view and manage a meta volume. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Name** — Volume name.
- ◆ **Volume Identifier** — Volume identifier.
- ◆ **Type** — Volume configuration.
- ◆ **Status** — Volume status.
- ◆ **Reserved** — Whether the volume is reserved.
- ◆ **Capacity (GB)** —Volume capacity in GBs.
- ◆ **Capacity (Cylinders)** —Volume capacity in cylinders.
- ◆ **Emulation** — Volume emulation.
- ◆ **Symmetrix ID** — Symmetrix system on which the volume resides.
- ◆ **Symmetrix Volume ID** — Symmetrix volume name/number.
- ◆ **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- ◆ **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- ◆ **Nice Name** — Nice name generated by Symmetrix Engineuity.
- ◆ **WWN** — World Wide Name of the volume.
- ◆ **DG Name** — Name of the device group in which the volume resides, if applicable.
- ◆ **CG Name** — Name of the device group in which the volume resides, if applicable.
- ◆ **Attached BCV** — Defines the attached BCV to be paired with the standard volume.

- ◆ **Attached VDEV TGT Volume** — Volume to which this source volume would be paired.
- ◆ **RDF Type** — RDF configuration.
- ◆ **Geometry - Type** — Method used to define the volume's geometry.
- ◆ **Geometry - Sectors per Track** — Number of sectors per track, as defined by the volume's geometry.
- ◆ **Geometry - Tracks per Cylinder** — Number of tracks per cylinder, as defined by the volume's geometry.
- ◆ **Geometry - 512 Block Bytes** — Number of 512 blocks, as defined by the volume's geometry.
- ◆ **SSID** — Subsystem ID.
- ◆ **Capacity (Tracks)** — Capacity in tracks.
- ◆ **SA Status** — Volume SA status.
- ◆ **Host Access Mode** — Host access mode.
- ◆ **Pinned** — Whether the volume is pinned.
- ◆ **RecoverPoint Tagged** — Whether the volume is tagged for RecoverPoint use.
- ◆ **Service State** — Service state.
- ◆ **Defined Label Type** — Type of user-defined label.
- ◆ **Dynamic RDF Capability** — RDF capability of the volume.
- ◆ **Mirror Set Type** — Mirror set for the volume and the volume characteristic of the mirror.
- ◆ **Mirror Set DA Status** — Volume status information for each member in the mirror set.
- ◆ **Mirror Set Invalid Tracks** — Number of invalid tracks for each mirror in the mirror set.
- ◆ **Priority QoS** — Priority value assigned to the volume. Valid values are 1 (highest) through 16 (the lowest).
- ◆ **Dynamic Cache Partition Name** — Name of the cache partition.

The following controls are available:

- ◆ **Create** — See [Managing Meta Volumes](#) on page 213.
- ◆ **Add Member** — See [Adding meta members](#) on page 215.
- ◆ **Dissolve** — See [Dissolving meta volumes](#) on page 216.
- ◆ **Convert** — See [Converting meta volumes](#) on page 216.
- ◆ **VLUN Migration** — See [Migrating regular volumes](#) on page 385.
- ◆ **Unpin** — See [Pinning and unpinning volumes](#) on page 191.
- ◆ **Pin** — See [Pinning and unpinning volumes](#) on page 191.

- ◆ **Assign Symmetrix Priority** — See [Assigning Symmetrix priority](#) on page 205.
- ◆ **Assign Dynamic Cache Partition** — See [Assigning volumes to dynamic cache partitions](#) on page 131.
- ◆ **Set Volume Status** — See [Setting volume status](#) on page 210.
- ◆ **Set Volume Attributes** — See [Setting volume attributes](#) on page 211.
- ◆ **Change Volume Configuration** — See [Changing volume configuration](#) on page 205.
- ◆ **Unmap** — See [Unmapping volumes](#) on page 208.
- ◆ **Map** — See [Mapping volumes](#) on page 206.
- ◆ **Untag for RecoverPoint** — See [Tagging and untagging volumes for RecoverPoint](#) on page 410.
- ◆ **Tag for RecoverPoint** — See [Tagging and untagging volumes for RecoverPoint](#) on page 410.

Related Objects panel

The **Related Objects** panel links you to views displaying objects associated with the meta volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Meta Members - 2** will open a view listing the two members in the meta volume.

Disk groups

Renaming disk groups

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group and click **View Details** to open its **Details** view.
4. Click **Rename** to open the **Rename Disk Group** dialog box.
5. Type the **New Name** and click **Run Now**.

Removing disks from disk groups

This procedure explains how to remove disk from disk groups.

Before you begin:

Only disks in external disk groups can be removed from a group.

To remove disks from a disk group:

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group from the list and click **View Details** to open its **Details** view.
4. From the **Related Objects** panel, select **Disks** to open the **Disks for Disk Group** view.
5. Select a disk from the list and click **Remove**.
6. Click **Remove** in the confirmation message.

Deleting disk groups

This procedure explains how to delete disk groups.

Before you begin:

Only empty external disk groups can be deleted.

To delete disk groups:

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Groups** list view.
3. Select one or more disk groups and click **Delete**.
4. Click **Delete** in the confirmation message.

Viewing disk groups

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Groups** list view.

Use this list view to view and manage disk groups.

The following properties display:

- ◆ **Name** — Name of disk group name, format: *number -- name*.
- ◆ **Technology** — Technology type for the disk group.
- ◆ **Disks** — Number of disks in the disk group.
- ◆ **Used Capacity (%)** — Percent total used capacity of the disk group, displayed in bar graph format and the actual percent number.
- ◆ **Total Capacity (GB)** — Total capacity in GB of the disk group.
- ◆ **Disk Location** — Indicates whether disk is internal or external.

The following controls are available:

- ◆ **View Details** — See [Viewing disk group details](#) on the facing page.
- ◆ **Delete** — See [Deleting disk groups](#) on previous page.

Viewing disk group details

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group from the list and click **View Details** to open the disk group **Details** view.

The disk group **Details** view allows you to view and manage disk groups. This view contains [Properties](#), [Related Objects](#), [Performance Views](#), and [Graphs](#) panels.

Properties panel

The following properties display:

- ◆ **Name** — User-defined disk group name, format: *number -- name*.
- ◆ **Technology** — Technology type for the disk group.
- ◆ **Number of Disks** — Number of disks in the disk group.
- ◆ **Used Capacity (GB)** — Total used capacity in GB of the disk group.
- ◆ **Free Capacity (GB)** — Total free capacity in GB of the disk group.
- ◆ **Total Capacity (GB)** — Total capacity in GB of the disk group.

The following controls are available:

- ◆ **Rename** — See [Renaming disk groups](#) on page 221.
- ◆ **Delete** — See [Deleting disk groups](#) on page 221.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the disk group. Each group link is followed the name of the group, or by a number, indicating the number of objects in the corresponding view. For example, clicking **Disks - 2** opens the view listing the two disks contained in the disk group.

Performance panel

The **Performance** panel links you to the performance analyze views for the group.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Graph panel

The **Graph** panel provides a graphic representation of the percentage of disk group used capacity to the disk group total capacity.

Viewing disks in disk group

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Disks** to open the **Disks for Disk Group** list view.
Use **Disks for Disk Group** list view to view and manage data disks in the disk group.

The following properties display:

- ◆ **Spindle** — Disk Spindle ID.
- ◆ **Dir** — Disk director ID.
- ◆ **Int** — DA SCSI path.
- ◆ **TID** — Disk SCSI ID.
- ◆ **Vendor ID** — Disk vendor.
- ◆ **Product Revision** — Product version number.
- ◆ **Hypers** — Number of disk hypers.
- ◆ **Total Cap (GB)** — Disk capacity.
- ◆ **Used (%)** — Percent of disk capacity.

The following controls are available:

- ◆ **View Details** — See [Viewing disk details](#) below.
- ◆ **Remove** — See [Removing disks from disk groups](#) on page 221.

Viewing disk details

1. Select the Symmetrix system.
2. Select **Storage > Disk Groups** to open the **Disk Group** list view.
3. Select the disk group from the list and click **View Details** to open the **Details** view.
4. In the **Related Objects** panel, select **Disks** to open the **Disks for Disk Group** list view.
5. From this list view, select a disk and click **View Details** to open the details view for the disk.

This view allows you to view the data disk details. This view contains three panels, [Properties](#), [Related Objects](#), and [Graphs](#).

Properties panel

The following properties display:

- ◆ **Spindle** — Spindle ID.
- ◆ **Disk ID** — Disk Identification.
- ◆ **Int** — DA SCSI path.

- ◆ **TID** — Disk SCSI ID.
- ◆ **External WWN** — World Wide Name of the external LUN.
- ◆ **Disk Group** — Disk group number.
- ◆ **Disk Location** — Location of disk.
- ◆ **Disk Technology** — Disk technology type.
- ◆ **Speed RPM** — Physical disk revolutions per minute.
- ◆ **Form Factor** — Form factor of the disk.
- ◆ **Vendor ID** — Disk vendor ID.
- ◆ **Product Revision** — Product revision number.
- ◆ **Serial ID** — Serial number.
- ◆ **Disk Blocks** — Number of disk blocks.
- ◆ **Actual Disk Blocks** — Actual number of disk blocks.
- ◆ **Block Size** — Size of each block.
- ◆ **Cap (GB)** — Useable disk capacity in Gigabytes.
- ◆ **Used (%)** — Percentage of used disk capacity to the total disk capacity.
- ◆ **Free Cap (GB)** — Free disk capacity in Gigabytes.
- ◆ **Actual Cap (GB)** — Actual disk capacity in Gigabytes
- ◆ **Rated Disk Capacity (GB)** — Rated capacity of the disk.
- ◆ **Spare Disk** — Indicates if disk is a spare.
- ◆ **Encapsulated** — If the disk is external, this indicates if it is encapsulated (True) or not (False).
- ◆ **Disk Service State** — Indicates disk service state.

The following control is available:

- ◆ **Remove** — See [Removing disks from disk groups](#) on page 221.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the disk group. Each group link is followed the name of the group, or by a number, indicating the number of objects in the corresponding view. For example, clicking **Hypers - 2** opens the view listing the two hypers contained in the disk.

Graph panel

The **Graph** panel provides a graphic representation of the percentage of used disk capacity to the total disk capacity.

Virtual Provisioning

DATA volumes

Creating DATA volumes

This procedure explains how to create DATA volumes from the **Volume Dashboard**. You can also create them by clicking the **Create volumes** option in the **Common Tasks** panel.

Before you begin:

DATA volumes are used with Symmetrix Virtual Provisioning, which requires Engenuity 5773 or higher.

To create DATA volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Private**.
6. Select **DATA** as the **Configuration**.
7. Select the **Disk Technology**.
External disk technology is an option if the Symmetrix system has FTS (Federated Tiered Storage) enabled and available external storage.
8. Select the **Emulation** type.
9. Select the **RAID Protection** level.
10. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
11. To add the new volumes to a specific thin pool, select one from **Add to Pool**. Pools listed are filtered on technology, emulation, and protection type.
12. Click one of the following:
 - ◆ **Add to Job List** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. Select the **Disk Group** (number and name) in which to create the volumes. The list of disk groups is already filtered based on technology type selected above.
2. To enable the new volumes in the pool, select **Enable volume in pool**.
3. Click either of the following:

- ◆ **Add to Job List** to create the volumes at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
- ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Activating and deactivating DATA volumes

This procedure explains how to activate or deactivate DATA volumes in a thin pool. Activating volumes is essentially the same thing as enabling volumes; however, the activate operation is not allowed if draining is in progress.

After activation, the volumes will go into the Enabled state.

Before you begin:

You can only activate deactivated DATA volumes with used tracks.

To activate/deactivate DATA volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select **DATA**. the **DATA Volumes** list view opens.
4. To activate volumes, right-click the volumes and select **Activate**. To deactivate volumes, right-click the volumes and select **Deactivate**.
5. Click **OK** in the confirmation message.

Enabling and disabling DATA volumes

This procedure explains how to enable or disable DATA volumes for use in a pool. The volumes in the pool do not all have to be in the same state (enabled or disabled). If all the volumes in a pool are disabled, then the pool is disabled. If at least one volume in a pool is enabled, then the pool is enabled.

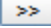
Before you begin:

To disable a volume, all sessions must be terminated, and have no used tracks.

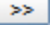
To enable/disable DATA volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select **DATA**. the **DATA Volumes** list view opens.
4. To enable volumes, right-click them and select **Enable**. To disable volumes, right-click them and select **Disable**.
5. Click **OK** in the confirmation message.

Start draining DATA volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select **DATA**.
4. Click **View** to open the **DATA Volumes** list view.
5. Select one or more volumes, click more  and select **Start Draining**.
6. Click **OK** in the confirmation message.

Stop draining DATA volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select **DATA**.
4. Click **View** to open the **DATA Volumes** list view.
5. Select one or more volumes, click more  , and select **Stop Draining**.
6. Click **OK** in the confirmation message.

Viewing DATA volumes

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select **DATA**.
4. Click **View** to open the **DATA Volumes** list view.
5. You can also view and manage DATA volumes from **Storage > Thin Pools > *thin pool name* > DATA Volumes**.

Use this list view to view and manage the volumes.

The following properties display:

- ◆ **Name** — Volume name.
- ◆ **Type** — Volume configuration.
- ◆ **Pool Name** — Pool in which the volume resides.
- ◆ **Pool Type** — Type of pool in which the volume resides.
- ◆ **Status** — Volume status.
- ◆ **Reserved** — Whether the volume is reserved.
- ◆ **Capacity** — Volume capacity.
- ◆ **Emulation** — Volume emulation.
- ◆ **State** — Whether the volume is enabled or disabled in the pool.
- ◆ **Session Status** — Whether volume is active or inactive.

The following controls are available:

- ◆ **Create** — See [Virtual Provisioning](#) on page 226.
- ◆ **View Details** — See [Viewing DATA volume details](#) below.
- ◆ **Delete** — See [Deleting volumes](#) on page 203.
- ◆ **Enable** — See [Enabling and disabling DATA volumes](#) on page 227.
- ◆ **Disable** — See [Enabling and disabling DATA volumes](#) on page 227.
- ◆ **Activate** — See [Activating and deactivating DATA volumes](#) on page 227.
- ◆ **Deactivate** — See [Activating and deactivating DATA volumes](#) on page 227.
- ◆ **Start Draining** — See [Start draining DATA volumes](#) on previous page.
- ◆ **Stop Draining** — See [Stop draining DATA volumes](#) on previous page.
- ◆ **Assign Dynamic Cache Partition** — [Assigning volumes to dynamic cache partitions](#) on page 131
- ◆ **Duplicate Volume** — [Duplicating volumes](#) on page 203

Viewing DATA volume details

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, expand the **Private Volume** folder and select **DATA**.
4. Click **View** to open the **DATA Volumes** list view.
5. Select a DATA volume and click **View Details** to open its **Details** view.

The **Details** view allows you to view and manage a volume. This view contains two panels: [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Name** — Volume name.
- ◆ **Volume Identifier** — Volume identifier.
- ◆ **Type** — Volume configuration.
- ◆ **Status** — Volume status.
- ◆ **Encapsulated Volume** — Whether external volume is encapsulated. Relevant for external disks only.
- ◆ **Encapsulated WWN** — World Wide Name for encapsulated volume. Relevant for external disks only.
- ◆ **Reserved** — Whether the volume is reserved.
- ◆ **Capacity (GB)** — Volume capacity in GBs.

- ◆ **Capacity (Cylinders)** —Volume capacity in cylinders.
- ◆ **Emulation** — Volume emulation.
- ◆ **Symmetrix ID** — Symmetrix system on which the volume resides.
- ◆ **Symmetrix Volume ID** — Symmetrix volume name/number.
- ◆ **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- ◆ **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- ◆ **Nice Name** — Nice name generated by Symmetrix Enginuity.
- ◆ **WWN** — World Wide Name of the volume.
- ◆ **DG Name** — Name of the device group in which the volume resides, if applicable.
- ◆ **CG Name** — Name of the device group in which the volume resides, if applicable.
- ◆ **Attached BCV** — Defines the attached BCV to be paired with the standard volume.
- ◆ **Attached VDEV TGT Volume** — Volume to which this source volume would be paired.
- ◆ **RDF Type** — RDF configuration.
- ◆ **Geometry - Type** — Method used to define the volume's geometry.
- ◆ **Geometry - Sectors per Track** — Number of sectors per track, as defined by the volume's geometry.
- ◆ **Geometry - Tracks per Cylinder** — Number of tracks per cylinder, as defined by the volume's geometry.
- ◆ **Geometry - 512 Block Bytes** — Number of 512 blocks, as defined by the volume's geometry.
- ◆ **Geometry - Limited** — Indicates whether the volume is geometry limited.
- ◆ **SSID** — Subsystem ID.
- ◆ **Capacity (Tracks)** — Capacity in tracks.
- ◆ **SA Status** — Volume SA status.
- ◆ **Host Access Mode** — Host access mode.
- ◆ **Pinned** —Whether the volume is pinned.
- ◆ **Service State** — Service state.
- ◆ **Defined Label Type** — Type of user-defined label.
- ◆ **Dynamic RDF Capability** — RDF capability of the volume.
- ◆ **Mirror Set Type** — Mirror set for the volume and the volume characteristic of themirror.

- ◆ **Mirror Set DA Status** — Volume status information for each member in the mirror set.
- ◆ **Mirror Set Invalid Tracks** — Number of invalid tracks for each mirror in the mirror set.
- ◆ **Priority QoS** — Priority value assigned to the volume. Valid values are 1 (highest) through 16 (the lowest).
- ◆ **Dynamic Cache Partition Name** — Name of the cache partition.

The following controls are available:

- ◆ **Create** — [Virtual Provisioning](#) (page 226).
- ◆ **Delete** — [Deleting volumes](#) (page 203)
- ◆ **Enable** — [Enabling and disabling DATA volumes](#) on page 227
- ◆ **Disable** — [Enabling and disabling DATA volumes](#) on page 227
- ◆ **Activate** — See [Activating and deactivating DATA volumes](#) on page 227.
- ◆ **Deactivate** — See [Activating and deactivating DATA volumes](#) on page 227.
- ◆ **Stop Draining** — See [Stop draining DATA volumes](#) on page 228.
- ◆ **Start Draining** — See [Start draining DATA volumes](#) on page 228.
- ◆ **Assign Dynamic Cache Partition** — [Assigning volumes to dynamic cache partitions](#) on page 131
- ◆ **Duplicate Volume** — [Duplicating volumes](#) on page 203

Related Objects panel

The **Related Objects** panel links you to views displaying objects contained in and associated with the DATA volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Back End Paths- 2** will open a view listing the two directors associated with the volume.

Thin pools

Creating thin pools

When creating thin pools, Unisphere for VMAX works on a best effort basis, 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.

To create a thin pool:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.

3. Click **Create** to open the **Create Thin Pool** dialog box. When this dialog box first opens, the chart displays the configured and unconfigured space on the selected Symmetrix system. Once you select a disk technology later in this procedure, and therefore a disk group, this chart will display the configured and unconfigured space of the selected group.
4. Type the **Pool Name**. Thin pool names can contain up to 12 alpha-numeric characters. The only special character allowed is the underscore (_); however, the name cannot start or end with an underscore.
5. Select the **Disk Technology** on which the pool will reside.
6. Select the RAID **Protection** level for the DATA volumes to use in the pool.
7. Select an **Emulation** type for the pool.
8. Specify the capacity that the template will find or create by doing one of the following:
 - ◆ Typing the **Number of Volumes** and selecting the capacity for each volume (**Volume Capacity**).
 - ◆ Selecting the **Volume Capacity**.
9. Click one of the following:
 - ◆ **OK** to review your selections and ultimately create the pool.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. Select the **Disk Group** containing the DATA volumes to use in the pool.
2. To specify the percentage of the pool's capacity to enable, select **Enable Max Subscription** and type a percentage.
3. Type the **Rebalancing Variance**. This is the target volume utilization variance for the rebalancing algorithm. The rebalancing algorithm attempts to level data distribution in a pool so that the percentage utilization of any volume in the pool is within the target variance of the percentage utilization of any other volume in the pool. Possible values range from 1 to 50%, with the default value being 1%. This field is only available when creating a thin pool on a Symmetrix system running Enginuity 5875 or higher.
4. Type the **Maximum Rebalancing Scan Device Range**. This is the maximum number of volumes in the pool on which the rebalancing algorithm will concurrently operate. Possible values range from 1 to 1024, with the default value being 256. This field is only available when creating a thin pool on a Symmetrix system running Enginuity 5875 or higher.
5. To specify the percentage of the capacity of the thin pool that will be reserved for non-FAST activities, select **Pool Reserved Capacity** and type a value. If the free space in the pool (as a percentage of pool-enabled capacity) falls below this value, the FAST controller does not move any more chunks into the pool. Specifying a value here will override the system-wide PRC value. Possible values range from 1 to 8.
6. To enable the DATA volumes in the pool for use, select **Enable DATA Volume for Use**.
7. To enable FAST VP compression for the volumes in a thin pool, select **Enable VP Compression**. This feature maximizes the storage capacity usage within the pool by compressing its volumes.
8. Click **OK**.
9. Verify your selections in the **Create Thin Pool - Summary** page, and click one of the following:

- ◆ **Add to Job List** to create the pool at a later time, as described in [Scheduling jobs](#) on page 100, or from the job list as described in [Viewing the job list](#) on page 102.
- ◆ Expand **Add to Job List** and click **Run Now** to create the pool now.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Expanding thin pools

Expanding thin pools refers to the process of increasing the amount of pool storage accessible to a thin volume by either adding a predefined capacity to the pool, or by increasing the pool's capacity by a percentage.

To expand a thin pool:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **Expand a thin pool** to open the **Expand Thin Pool** dialog box. The chart on this dialog box displays the configured and unconfigured space of the disk group containing the pool's DATA volumes.
4. Select how to expand the pool, either by **Capacity** or **Percentage**:
 - ◆ **Capacity** — The **Volume Capacity** field defaults to the first data volume size in the pool. All volume sizes contained in the pool are available. Type the **Extra Pool Capacity** in GBs.
 - ◆ **Percentage** — Type an amount in the **Percentage Increase** field.
5. *Optional:* Click **Show Advanced** to enable **Write Balancing**.
6. Click **OK**.
7. Verify your selections in the **Expand Thin Pool - Summary** page, and click **Add to Job List** or **Run Now**.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Draining thin pools



This procedure explains how to re-balance data across all the DATA volumes in a thin pool. This procedure is typically performed after expanding a thin pool.

Before you begin:

- ◆ The drain operation is not supported with any ongoing replication operation.
- ◆ You can only drain deactivated DATA volumes. For instructions, refer to [Activating and deactivating DATA volumes](#) on page 227.

- ◆ The drain must not cause the enabled volumes to end up with greater than 90% utilization in the pool. To calculate this, EMC® Unisphere for VMAX adds the total used tracks on the enabled volumes and the total used tracks on the volumes that will be drained and divides this sum by the total number of tracks on all the enabled volumes. If the result is greater than 90% the drain request is blocked.
- ◆ The number of volumes that are draining at any time are limited to 20% total of the number of volumes to drain (or draining) plus the number of enabled volumes. This limits the impact on the system.
- ◆ This feature is only supported on Symmetrix DMX systems running Enginuity 5773.150 or higher.

To drain thin pools:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **DATA Volumes** to open the **DATA Volumes for Thin Pool** list view.
5. Select one or more volumes, click more  , and select **Start Draining**.
6. Click **OK** in the confirmation message. This will put the volumes in a Draining state.
7. Monitor the draining until it reaches an acceptable percentage. This will require you to refresh the view. If you do not monitor the draining, eventually all data will be drained from the volumes and they will go into a Disable state.
8. When a volume reaches an acceptable level, select it, click more  , and select **Stop Draining**.
9. Click **OK** in the confirmation dialog. This will put the volume in an Enabled state.
10. If you are draining multiple devices, repeat steps 8-10 until all the volumes are drained to an acceptable percentage.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.



Starting and stopping thin pool write balancing

Write balancing thin pools refers to the process of rebalancing allocated capacity across all the DATA volumes in the pool. This procedure is typically performed after expanding a thin pool.

Before you begin:

- ◆ You can only perform this procedure on an enabled thin pool with at least one thin volume bound to it.
- ◆ While write balancing is going on, all pool operations can still occur.
- ◆ Write balancing requires Engenuity 5874.207.166 or higher.

To write balance a thin pool:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **View Details** to open the thin pool's details view.
4. Start write balancing by clicking more , and selecting **Start Write Balancing**.
5. Click **OK** in the confirmation message. This will put the pool in a Balancing state.
6. Monitor the balancing until it reaches an acceptable percentage.
7. Select the thin pool, click more , and select **Stop Write Balancing**.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Deleting thin pools

Before you begin:

You can only delete empty thin pools. For instructions, refer to Adding/Removing thin pool members.

To delete a thin pool:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **Delete**.
4. Click **Delete** in the confirmation message.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Adding/Removing thin pool members

This procedure explains how to add/remove members from a thin pool.

Before you begin:

Before you can remove a thin pool member (DATA volume), you must first [disable it](#).

To add/remove thin pool members:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **View Details** to open the thin pool's details view.
4. In **Related Objects**, click **DATA Volumes** to open the **DATA Volumes for Thin Pool** view.
5. Do the following, depending on whether you are adding or removing volumes from the pool:

Adding volumes:

- a. Click **Add Volumes to Pool** to open the **Add Volumes to Thin Pool** dialog box.
- b. Select one or more volumes and click **OK**.

Removing volumes:

- a. Select one or more volumes and click **Remove**.
- b. Click **OK** in the confirmation message.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Enabling and disabling thin pool members

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **View Details** to open the thin pool's details view.
4. In **Related Objects**, click **DATA Volumes** to open the **DATA Volumes for Thin Pool** view.
5. To enable members, select them and click **Enable**.
To disable members, select them and click **Disable**.
6. Click **OK** in the confirmation message.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Starting and stopping thin pool allocation


The following describes how to start and stop allocating thin pool capacity for thin volumes. The thin volumes can be individual volumes or volumes in a device group or storage group.

Before you begin:

- ◆ You can only allocate thin pool capacity to bound thin volumes.
- ◆ Thin pool allocation for individual thin volumes requires Engenuity 5874 or higher.
- ◆ Thin pool allocation for thin volumes in a device group or storage group requires Engenuity 5875 or higher.

- ◆ When starting or stopping thin pool allocation for a device group, the operation is limited to the STD volumes in the group.

To start thin pool allocation:

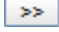
1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes Dashboard**.
3. In the **Volume Type** panel, select the type of volume on which to start the allocation process.
4. Click **View** to open the volume list view.
5. Select one or more volumes, click more , and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
6. To view configuration, capacity, and emulation information for the selected volumes, click **Show Selected Volumes**.
7. Select **Start Allocate**. In addition, on Symmetrix systems running Enginuity 5875 Q2 20011 SR, you can optionally specify to persist preallocated capacity on the thin volumes by selecting the **Persist preallocated capacity through reclaim or copy** option. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations.
8. To allocate pool capacity for entire volumes, select **Full Device**.
9. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To allocate pool capacity for a region of the thin volumes, type a starting and ending cylinder number.
2. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in [Symmetrix Preferences](#) for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
3. Click either of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

To stop thin pool allocation:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes Dashboard**
3. In the **Volume Type** panel, select the type of volume on which to stop the allocation process.

4. Click **View** to open the volume list view.
5. Select one or more volumes, click more , and select **Stop Allocate/Free** to open the **Stop Allocate/Free** dialog box.
6. To view configuration, capacity, and emulation information for the selected volumes, click **Show Selected Volumes**.
7. Select **Stop Allocate**.
8. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
9. Click **Run Now**.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.


Starting and stopping space reclamation

The following describes how to manage the space reclamation options, freeing unused pool capacity and reclaiming unwritten tracks from thin volumes.

Before you begin:

- ◆ You can only reclaim thin pool capacity from bound thin volumes.
- ◆ Thin pool reclamation for individual thin volumes requires Engenuity 5874 or higher.
- ◆ Thin pool reclamation for thin volumes in a device group or storage group requires Engenuity 5875 or higher.
- ◆ When starting or stopping thin pool reclamation for a device group, the operation is limited to the STD volumes in the group.

To start space reclamation:


1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes Dashboard**.
3. In the **Volume Type** panel, expand the **Virtual Volume** folder and select the type of volume on which to start the reclamation process.
4. Click **View** to open the volume list view.
5. Select one or more volumes, click more , and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
6. To view configuration, capacity, and emulation information for the selected volumes, click **Show Selected Volumes**.
7. To start freeing unused pool capacity, select **Start Free**.
8. To start reclaiming unwritten tracks and those written completely with zeros, select **Start Reclaim**. This option is only available on Symmetrix systems running Engenuity 5875 Q2 2011 SR or higher.

9. To reclaim tracks that are unwritten or zero-based, even if they are marked as persistent, select **Reclaim persistent capacity**.
10. To perform the operation on the entire volume, select **Full Volume**.
11. Click one of the following:
 - ◆ **Run Now** to start the task now.
 - ◆ **Add to Job List** to schedule it for a later time.
 - ◆ **Show Advanced** to continue setting the advanced options, as described in setting advanced options next.

Setting Advanced options:

1. To perform the operation on an allocated region of the thin volumes, type a starting and ending cylinder number.
2. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
3. Click one of the following:
 - ◆ **Run Now** to start the task now.
 - ◆ **Add to Job List** to start the task at a later time.

To stop space reclamation:

1. From the **System Selector**, select the Symmetrix system containing the thin volumes.
2. In the navigation bar, click **Storage** to open the **Storage** section.
3. Click **Volumes** to open the **Volumes Dashboard**.
4. In the **Volume Type** panel, expand the **Virtual Volume** folder and select the type of volume on which to stop the reclamation process.
5. Click **View** to open the volume list view.
6. Select one or more volumes, click more  and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
7. To view configuration, capacity, and emulation information for the selected volumes, click **Show Selected Volumes**.
8. Select **Stop Reclaim**. In addition, on Symmetrix systems running Enginuity 5875 Q2 20011 SR, you can optionally specify to free tracks that are unwritten or zero-based, even if they are marked persistent. This option is only available on Symmetrix systems running Enginuity 5875 Q2 2011 SR.
9. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**. The default values for **Reserve** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
10. Click either of the following:
 - ◆ **Run Now** to start the task now.

- ◆ **Add to Job List** to schedule it for a later time, as described in [Scheduling jobs](#) on page 100.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Viewing thin pools

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.

The **Thin Pools** list view allows you to view and manage thin pools on a Symmetrix system.

The following properties display:

- ◆ **Name** — Name of the thin pool.
- ◆ **Technology** — Disk technology on which the pool resides.
- ◆ **Disk Location** — Whether the disk group is **Internal** to the Symmetrix system or an **External** Symmetrix system or storage device.
- ◆ **Configuration** — Configuration of the pool.
- ◆ **Allocated Capacity** — Percentage of the pool that is allocated.
- ◆ **Capacity** — Capacity of the pool in GB.

The following controls are available:

- ◆ **Create** — See [Thin pools](#) on page 231.
- ◆ **Expand** — See [Expanding thin pools](#) on page 233.
- ◆ **View Details** — See [Viewing thin pool details](#) below.
- ◆ **Delete** — See [Deleting thin pools](#) on page 235.

Viewing thin pool details

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the pool and click **View Details** to open its **Details** view.

The thin pool **Details** view allows you to display and manage a thin pool. It contains [Properties](#), [Related Objects](#), [Performance](#), and [Graphs](#) panels.

Properties panel

The following properties display:

- ◆ **Name** — Name of the pool. To rename a pool, type a new name over the existing and click **Apply**. Thin pool names can contain up to 12 alpha-numeric characters. The only special character allowed is the underscore (`_`); however, the name cannot start or end with an underscore.
- ◆ **RAID Protection** — RAID protection level for the DATA volumes in the pool.
- ◆ **Technology** — Disk technology on which the pool resides.

- ◆ **Emulation** — Emulation type for the pool.
- ◆ **Total Capacity** — Total capacity of the pool.
- ◆ **Free Capacity** — Free capacity in the pool.
- ◆ **Thin Volumes** — Number of thin volumes bound to the pool.
- ◆ **Enabled Volumes** — Number of enabled DATA volumes in the pool.
- ◆ **Disabled Volumes** — Number of disabled DATA volumes in the pool.
- ◆ **% Allocated** — Percent of pool used.
- ◆ **Maximum Subscription** — Acceptable oversubscription ratio for the pool.
- ◆ **% Subscription** — Current subscription percentage.
- ◆ **Enabled Capacity** — Sum of capacity of all enabled DATA volumes in the pool.
- ◆ **Allocated Capacity** — Pool capacity allocated to thin volumes.
- ◆ **State** — Pool state (Enabled, Disable, Balancing).
- ◆ **Rebalance Variance** — Target volume utilization variance for the rebalancing algorithm. The rebalancing algorithm attempts to level data distribution in a pool so that the percentage utilization of any volume in the pool is within the target variance of the percentage utilization of any other volume in the pool. Possible values range from 1 to 50%, with the default value being 1%.
- ◆ **Maximum Volumes per Rebalance Scan** — Maximum number of volumes in the pool on which the rebalancing algorithm will concurrently operate. To change this number, type a new value over the existing and click **Apply**. Possible values range from 1 to 1024, with the default value being 256. This field only applies to thin pool on a Symmetrix system running Enginuity 5875 or higher.
- ◆ **Pool Egress Counter** — Number of track groups freed from the thin pool as a result of a FAST related data movement.
- ◆ **Pool Ingress Counter** — Number of track groups allocated in the thin pool as a result of a FAST related data movement.
- ◆ **Enable VP Pool Compression** — Enables (when selected) or disables (when cleared) VP pool compression for the pool. VP pool compression compresses data on the volumes in the pool to save space. VP Pool compression is available for thin pools only not Snap or DSE pool types. VP compression cannot be enabled on a thin pool of encapsulated volumes.
- ◆ **Total Compressed Capacity** — Amount of allocated pool capacity that is compressed.
- ◆ **Pool Compression Ratio** — Percentage of allocated pool capacity that is compressed.

The following controls are available:

- ◆ **Create** — See [Thin pools](#) on page 231.
- ◆ **Expand** — See [Expanding thin pools](#) on page 233.
- ◆ **Bind** — See [Binding/Unbinding/Rebinding thin volumes](#) on page 244.
- ◆ **Assign Dynamic Cache Partition** — See [Assigning volumes to dynamic cache partitions](#) on page 131.
- ◆ **Stop Write Balancing** — See [Starting and stopping thin pool write balancing](#) on page 235.
- ◆ **Start Write Balancing** — See [Starting and stopping thin pool write balancing](#) on page 235.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the thin pool.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Objects panel

The **Related Objects** panel links you to views displaying objects contained in and associated with the thin pool. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **DATA Volumes - 2** will open a view listing the two DATA volumes in the pool.

Performance panel

The **Performance** panel links you to the performance monitor and analyze views for the thin pool.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Graphs panel

The **Graphs** panel provides a graphic representation of the thin pool's allocation as a percentage.

Thin volumes

Creating thin volumes

This procedure explains how to create thin volumes from the **Volume Dashboard**. You can also create them by clicking the **Create Volumes** option in the **Common Tasks** panel.

To create thin volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. Select a volume type and click **View** to open the volume list view.
4. Click **Create** to open the **Create Volume** dialog box.
5. In **Volume Type**, click **Virtual**.
6. Select **Configuration (TDEV or BCV + TDEV)** or thin volumes.
7. Select the **Emulation** type.

8. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
9. To bind the new volumes to a specific thin pool, select one from **Bind to Pool**. Only thin pools with enabled DATA volumes and matching emulation are available for binding (except AS/400 which will bind to an FBA pool).
10. Click one of the following:
 - ◆ **Add to Job List**. to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. If creating a thin volume or a thin BCV, you can specify the amount of space to preallocate on each volume, by selecting **Preallocate capacity for each new volume** selecting either **Full Volume** or **By Capacity**, and then typing a value. The amount of space must be less than or equal to the available space in the pool and less than or equal to the allocated size of the thin volume. This option is not available when creating VDEVs.
2. If you specified to preallocate space in the previous step, you can mark the preallocation on the thin volume as persistent by selecting **Persist preallocated capacity through reclaim or copy**. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations. This option is not available when creating VDEVs.
3. To assign **Dynamic Capability** to the volumes, select one of the following; otherwise, leave this field set to **None**.
 - ◆ **RDF1_Capable** — Creates a dynamic R1 RDF volume.
 - ◆ **RDF2_Capable** — Creates a dynamic R2 RDF volume.
 - ◆ **RDF1_OR_RDF2_Capable** — Creates a dynamic R1 or R2 RDF volume.
4. View **Enable SCSI3 Persistent Reservation** status — For Enginuity 5875 and higher this feature is pre-set by SYMAPI and cannot be changed. It is displayed as enabled for Enginuity 5875 and higher, except for CDK and AS/400 emulations.
If Auto meta is enabled on the system then it displays as enabled with a green check mark.
5. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name**:
 - ◆ **Name Only** — All volumes will have the same name.
 - ◆ **Name + VolumeID** — All volumes will have the same name with a unique Symmetrix volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
 - ◆ **Name + Append Number** — All volumes will have the same name with a unique decimal suffix appended to them. The suffix will start with the value specified for the **Append Number** and increment by 1 for each additional volume. Valid **Append Numbers** must be from 0 to 1000000. When using this option, the maximum number of characters allowed is 50.

For more information on naming volumes, refer to [Setting volume names](#) on page 212

6. Click one of the following:
 - ◆ **Add to Job List.** to create the volumes at a later time, as described in [Scheduling jobs](#) (page 100), or from the job list as described in [Viewing the job list](#) (page 102).
 - ◆ Expand **Add to Job List** and click **Run Now** to create the volumes now.

Binding/Unbinding/Rebinding thin volumes

You can bind /unbind/rebind thin volumes at the [volume](#), [pool](#), or [storage group](#) level.

Before you begin:

- ◆ Only one bind, unbind, or rebind operation can be performed on the same volume in any one config session.
- ◆ To bind/unbind thin volumes from the pool level or volume level, requires Engenuity 5773 or higher.
- ◆ To bind/unbind thin volumes from the group level, requires Engenuity 5875 or higher.
- ◆ As an alternative to unmapping/unmasking a volume prior to unbinding, you can make the volume Not Ready.
- ◆ A thin volume cannot be unbound from a pool if any of the following are true:
 - Volume is mapped to a front-end port or is in the Ready state
 - Volume is masked by VCM
 - Volume has active snap sessions
 - Volume is held
 - Volume is a source or target of a clone (src or tgt) session
 - Volume is a metamember
 - Volume is a part of enabled RDF CG group
 - Volume is an RDF volume
- ◆ The following apply just to the rebind operation:
 - The thin volume has to be in the Bound state.
 - The new binding has to comply with the oversubscription ratio of the new pool. The entire size of the volume being rebound will be considered when calculating the oversubscription.
 - If volumes in a range, device group, or storage group are bound to different pools, then all the volumes will be rebound to the specified pool.
 - If a thin volume is part of a storage group that is under FAST management, the thin volume can only be bound to a pool in a tier that is part of the FAST policy associated with the storage group. Therefore, the volume can only be rebound to a pool that is within the policy.
 - If all the volumes that are being rebound are already bound to the destination pool, an error returns. If some volumes get bound to a pool different than what they are currently bound to, the operation will return a success status.

To bind/unbind/rebind thin volumes at the volume level:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, select the type of volume to bind/unbind.
4. Click **View** to open the volume list view.
5. Do the following, depending on whether you are binding or unbinding thin volumes:

Binding:

- a. Select one or more volumes and click **Bind** to open the **Bind Thin Volumes** dialog box.
- b. From **Pool Name**, select the thin pool with which to bind the volume.
- c. To **Preallocate capacity for each volume**, select the option and specify the amount of each volume to preallocate, either **By Percent** or **By Capacity**.
- d. If you selected to preallocate space in the previous step, you can mark the preallocation as persistent by selecting **Persist preallocated capacity through reclaim or copy**. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations.
- e. To reserve the volumes, select **Reserve** and enter **Comments** and an **Expiration** date.
- f. Click **OK**.

Unbinding:

- a. Select one or more volumes and click **Unbind** to open the **Unbind Thin Volumes** dialog box.
- b. To view additional information on the selected volumes, click **Show selected volumes**.
- c. To reserve the volumes, select **Reserve** and enter **Comments** and an **Expiration** date.
- d. Click **Run Now**.
- e. Click **OK** in the confirmation message.

Rebinding:

- a. Select one or more volumes and click **Rebind** to open the **Rebind Thin Volumes** dialog box.
- b. Select the **Pool Name** with which to rebind the volumes.
- c. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**.
- d. Click **OK**.

To bind/unbind/rebind thin volumes at the pool level:

1. Select the Symmetrix system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **View Details** to open the thin pool's details view.
4. Do the following, depending on whether you are binding or unbinding volumes:

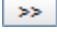
Binding:

- a. Click **Bind** to open the **Bind Thin Volumes** dialog box.
- b. Select one or more volumes.
- c. To **Preallocate capacity for each volume**, select the option and specify the amount of each volume to preallocate, either **By Percent** or **By Capacity**.
- d. If you selected to preallocate space in the previous step, you can mark the preallocation as persistent by selecting **Persist preallocated capacity through reclaim or copy**. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations.
- e. To reserve the volumes, select **Reserve** and enter **Comments** and an **Expiration** date.
- f. Click **OK**.

Unbinding:

- a. In **Related Objects**, click **Bound Volumes** to open the **Bound Volumes for Thin Pool** detail view.
- b. Select one or more volumes and click **Unbind**.
- c. Click **OK** in the confirmation message.

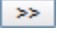
Rebinding:

- a. Select one or more volumes, click more , and select **Rebind** to open the **Rebind thin volumes** dialog box.
- b. Select the **Pool Name** with which to rebind the volumes.
- c. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**.
- d. Click **OK**.

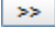
To bind/unbind/rebind thin volumes at the storage group level:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. Do the following, depending on whether you are binding or unbinding thin volumes:


Binding:

- a. Select the storage group, click more , and select **Bind SG** to open its **Bind Thin Volumes** dialog box.
- b. From **Pool Name**, select the thin pool with which to bind the volume.
- c. To **Preallocate capacity for each volume**, select the option and specify the amount of each volume to preallocate, either **By Percent** or **By Capacity**.
- d. If you selected to preallocate space in the previous step, you can mark the preallocation as persistent by selecting **Persist preallocated capacity through reclaim or copy**. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations.
- e. To reserve the volumes, select **Reserve** and enter **Comments** and an **Expiration** date.
- f. Click **OK**.

Unbinding:

- a. Select the storage group, click more  , and select **Unbind SG** to open its **Unbind Thin Volumes** dialog box.
- b. To view additional information on the selected volumes, click **Show selected volumes**.
- c. To reserve the volumes, select **Reserve** and enter **Comments** and an **Expiration** date.
- d. Click **Run Now**.
- e. Click **OK** in the confirmation message.

Rebinding:

- a. Select the storage group, click more  , and select **Rebind SG** to open the **Rebind Thin Volumes** dialog box.
- b. Select the **Pool Name** with which to rebind the volumes.
- c. To reserve the volumes, select **Reserve**. In addition you can also type reserve **Comments** and select an **Expiration Date**.
- d. Click **OK**.

For more information about thin pools and thin provisioning concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Enhanced Virtual LUN migration

Understanding Virtual LUN Migration

Virtual LUN Migration (VLUN Migration) enables transparent, nondisruptive data mobility for both disk group provisioned and virtually provisioned Symmetrix 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 Symmetrix system. Migrations are performed while providing constant data availability and protection.

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 Symmetrix 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 Symmetrix system volume.

Virtual LUN Migration can be used to migrate regular Symmetrix 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.


Migrating regular storage group volumes

This procedure explains how to migrate all the regular volumes in a storage group.

Before you begin:

- ◆ Virtual LUN migration requires Engenuity 5874 or higher.
- ◆ To migrate volumes when access control is enabled, you must grant Symmetrix Optimizer access rights to the volumes.

To migrate regular storage group volumes:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** view.
3. Select a storage group, click more , and select **VLUN Migration** to open the **LUN Migration** dialog box.
4. Type a **Migration session name**. Migration session names must be less than 32 characters long and are case sensitive.
5. Select a **Target Disk Group**.
6. Select the RAID **Protection type**.
7. Select **Target type**. Choose **Create new volumes** to migrate to unconfigured volumes or **Use existing volumes** to migrate to configured volumes.
8. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
9. Click **OK** to create the migration session.


Migrating regular volumes

This procedure explains how to migrate individual regular volumes.

Before you begin:

- ◆ To migrate volumes when access control is enabled, you must grant Symmetrix Optimizer access rights to the volumes.
- ◆ Virtual LUN migration requires Engenuity 5874 or higher.

To migrate regular volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** view.
3. Expand the **Regular Volumes** folder, select the desired **Volume Type**, and click **View** to open the volume list.
4. Select one or more volumes, click more , and select **VLUN Migration** to open the **VLUN Migration** dialog box.
5. Type a **Migration session name**. Migration session names must be less than 32 characters and are case sensitive.
6. Select a **Target Disk Group**.
7. Select the RAID **Protection type**.

8. Select the **Target type**. Choose **Create new volumes** to migrate to unconfigured volumes or **Use existing volumes** to migrate to configured volumes.
9. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
10. Click **OK**.

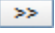
Migrating thin storage group volumes

This procedure explains how to migrate all the thin volumes in a storage group.

Before you begin:

- ◆ Virtual LUN migration requires Enginuity 5874 or higher.
- ◆ To migrate volumes when access controls is enabled, you must grant Symmetrix Optimizer access rights to the volumes.
- ◆ For instructions on migrating volumes on Symmetrix systems running Enginuity levels lower than 5874, see Migrating volumes.

To migrate thin storage group volumes:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** view.
3. Select a storage group, click more , and select **VLUN Migration** to open the **VLUN Migration** dialog box.
4. Type a **Migration session name**. Migration session names must be less than 32 characters and are case sensitive.
5. Select a **Target Pool**.
6. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
7. Click **OK**.


Migrating thin volumes

This procedure explains how to migrate individual thin volumes.

Before you begin:

- ◆ Virtual LUN migration requires Enginuity 5874 or higher.
- ◆ To migrate volumes when access controls is enabled, you must grant Symmetrix Optimizer access rights to the volumes.

To migrate selected thin volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** view.
3. Expand the **Virtual Volumes** folder, select **TDEV** Volume Type, and click **View** to open the **Thin Volumes** list.
4. Select one or more volumes, click more , and select **VLUN Migration** to open the **VLUN Migration** dialog box.
5. Type a **Migration session name**. The session name must be less than 32 characters long and is case sensitive.
6. Select a **Target Pool**.
7. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
8. Click **OK**.

Terminating a VLUN migration session

1. Select the Symmetrix system.
2. Select **Data Protection > Migration** to open the **Migration** sessions list view.
3. Select the migration session and click **Terminate** to open the **Terminate** confirmation dialog.
4. Click **OK**.

Viewing VLUN migration sessions

1. Select the Symmetrix system.
2. Select **Data Protection > Migration** to open the **Migration** sessions list view.

Use the this view to display and manage migration sessions.

The following properties display:

- ◆ **Name** — Migration session name.
- ◆ **Invalid Tracks** — Number of invalid tracks for the volume pair.
- ◆ **Status** — Migration session status.
- ◆ **Percentage** — Percentage of the session completed.

The following controls are available:

- ◆ **Terminate** — [Terminating a VLUN migration session](#) above
- ◆ **View Details** — [Viewing VLUN migration session details](#) on page 387

Storage templates

Understanding storage templates

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 expanding storage groups (Auto-provisioning Groups).
- ◆ Creating volumes (regular and thin).

In addition, storage templates can be shared (using import and export) between administrators and among Unisphere servers to help facilitate consistency across the storage environment.

Storage templates require Enginuity 5874 or higher and storage groups.

Creating storage templates

This procedure explains how to create storage templates.

Before you begin:

- ◆ Storage templates require Enginuity 5874 or higher.
- ◆ This feature is only available for a user with Admin or StorageAdmin permission.

To create a storage template:

1. Select **Storage > Storage Templates** to open the **Storage Template** list view.
2. Click **Create** to open the **Create Storage Template** wizard.
3. Type a **Storage Template Name**. Storage template names can be up to 64 character long.
4. Select the **Symmetrix ID** with which the template will be associated or select **ANY** to associate the template with any supported Symmetrix system.
5. Select whether to include **Virtual** or **Regular** volumes in the template.
6. Specify the capacity that the template will find or create capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**.
7. Click **Next**.

8. Do the following, depending on whether you are including virtual or regular volumes in the template:

Virtual volumes:

- a. Select a **Volume Configuration**.
- b. Select the **Thin Pool** to which you are going to bind the volumes.

Regular volumes:

- a. Select the **RAID Protection** level.
- b. If the template will be associated with a specific Symmetrix system, select the target **Disk Technology** and **Disk Group** (name and number).

9. Click one of the following:

- ◆ **Next** to verify your selection and create the template.
- ◆ **Show Advanced** to continue setting the advanced options as described next.

Setting Advanced options:

1. If you are including virtual volumes in the template, complete the following steps:
 - a. To specify how much of each volume to preallocate, select **Preallocated capacity for each volume**, select either **By Percent** or **By Capacity**, and type a value.
 - b. To maintain the amount of preallocated space through thin volume reclaims and volume copies, select **Persist preallocated capacity through reclaim or copy**.
2. To specify the **Method** to use when selecting volumes, select one of the following:
 - **Best Effort** — Specifies to satisfy as much as possible of the total requested capacity from existing volumes, and then create the volumes necessary to meet any shortfall.
 - **Existing Volumes** — Specifies to select from existing volumes.
 - **Create New Volumes** — Specifies to create new volumes.
3. To use SRDF volumes in the template, select one of the following; otherwise, leave this field set to **None**.
 - **None** — Specifies to not use SRDF volumes.
 - **RDF Dynamic** — Specifies to only use dynamic SRDF volumes.
 - **R1** — Specifies to only use R1 volumes. This option is only available when you are creating or modifying a template that uses the Find Existing method for locating volumes; except when the volumes are thin (TDEVs), in which case this option is not available.
4. Select the volume **Emulation**. The storage template uses only volumes of the specified emulation.
5. To include metavolumes in the template, select **Enable Meta** and complete the following steps:
 - a. Select the **Meta Config**. The storage template uses only meta volumes of the specified configuration.
 - b. Select the **Meta Member Capacity**. The storage template uses only meta volumes of the specified capacity.
6. Click **Next**.
7. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.

8. Click **Finish**.

Modifying storage templates

This procedure explains how to modify an existing storage template.

Before you begin:

- ◆ Storage templates require Enginuity 5874 or higher.
- ◆ The user must have Administrator or StorageAdmin permission.

To modify a storage template:

1. Select **Storage > Storage Templates** to open the **Storage Template** list view.
2. Select the template and click **Edit** to open the **Edit Storage Template** wizard.
3. Modify the template as you step through the wizard.
4. Click **Finish**.

Exporting storage templates

This procedure explains how to export storage templates to a temporary directory on the Unisphere server.

Before you begin:

- ◆ Storage templates require Enginuity 5874 or higher.
- ◆ The user must have Administrator or StorageAdmin permission.

To export storage templates:

1. Select **Storage > Storage Templates** to open the **Storage Template** list view.
2. Select one or more templates and click **Export**.
3. Click **OK** in the confirmation message.
4. Copy the templates from one Unisphere server to the other. Templates reside in the server's template directory (`Install_directory\SMAS\jboss\server\default-em\data\TemplateFiles`).

Importing storage templates

This procedure explains how to import storage templates to a Unisphere server.

Before you begin:

- ◆ Storage templates require Enginuity 5874 or higher.
- ◆ The user must have Administrator or StorageAdmin permission.
- ◆ One or more exported storage templates.

To import storage templates:

1. Select **Storage > Storage Templates** to open the **Storage Template** list view.
2. Click **Import** to open the **Import Storage Template** list.
All previously exported storage templates display.
3. Select one or more templates and click **OK**.

Deleting storage templates

This procedure explains how to delete storage templates.

Before you begin:

- ◆ Storage templates require Enginuity 5874 or higher.
- ◆ This feature is only available for a user with Administrator or StorageAdmin permission.

To delete a storage template:

1. Select **Storage > Storage Templates** to open the **Storage Template** list view.
2. Select the template and click **Delete**.
3. Click **Yes** in the confirmation box.

Viewing storage templates

1. Select the Symmetrix system.
2. Select **Storage > Storage Templates** to open the **Storage Template** list view.

The **Storage Template** list view allows you to view and manage storage templates.

The following properties display:

- ◆ **Storage Template** — Name of the template.
- ◆ **Symmetrix ID** — Symmetrix system associated with the template. Any indicates that the template is associated with any Symmetrix system running Enginuity 5874 or higher connected to the server.
- ◆ **Volume Type** — Type of volume specified in the template.
- ◆ **Protection/Volume Configuration** — RAID protection level/Thin or Thin BCV.

- ◆ **Total Request Capacity** — Capacity the template will find or create.
- ◆ **Emulation** — Emulation type of the disks specified in the template.
- ◆ **Method** — Method the template is using to select volumes. Possible values are:
 - **Best Effort** — Specifies to satisfy as much as possible of the total requested capacity from existing volumes, and then create the volumes necessary to meet any shortfall.
 - **Existing Volumes** — Specifies to select from existing volumes.
 - **Create New Volumes** — Specifies to create new volumes.

The following controls are available:

- ◆ **Create** — See [Creating storage templates](#) on page 252.
- ◆ **View Details** — See [Viewing storage template details](#) below.
- ◆ **Import** — See [Importing storage templates](#) on page 254.
- ◆ **Export** — See [Exporting storage templates](#) on page 254.
- ◆ **Delete** — See [Deleting storage templates](#) on previous page.

Viewing storage template details

1. Select the Symmetrix system.
2. Select **Storage > Storage Templates** to open the **Storage Template** list view.
3. Select the storage template and click **View Details** to open its **Details** view.

The storage template **Details** view allows you to view and manage storage templates.

The following properties display:

- ◆ **Template Name** — Name of the template. To rename a template, type a new name over the existing and click **Apply**. Storage template names can contain up to 64 characters.
- ◆ **Symmetrix ID** — The identifier of the associated Symmetrix system. Any indicates that the template is associated with all supported Symmetrix systems.
- ◆ **Volume Type** — The template specifies virtual or regular volumes.
- ◆ **Number of Volumes** — The number of volumes the template will find/create.
- ◆ **Volume Capacity** — Capacity of each volume.
- ◆ **Total Requested Capacity** — Total requested capacity.
- ◆ **Create Meta** — Whether the template will find/create meta volumes.
- ◆ **Meta Volume Configuration** — Configuration of the meta volumes (Striped or Concatenated).
- ◆ **Meta Member Capacity** — Capacity of each meta member.
- ◆ **Meta Member Count Including Head** — Number of meta members, including the head.

- ◆ **Protection/Volume Configuration** — Protection level/configuration of the volumes used in the template.
- ◆ **Disk Group** — Disk group from which the template is finding/creating volumes.
- ◆ **Thin Pool** — Thin pool to which the thin volumes are bound.
- ◆ **Preallocated Capacity for each volume** — How much of each volume is preallocated.
- ◆ **Persist Preallocation Through Reclaim or Copy** — Whether to maintain the amount of preallocated space through thin volume reclaims and volume copies.
- ◆ **Method** — Method used to select volumes. Possible values are:
 - **Best Effort** — Specifies to satisfy as much as possible of the total requested capacity from existing volumes, and then create the volumes necessary to meet any shortfall.
 - **Existing Volumes** — Specifies to select from existing volumes.
 - **Create New Volumes** — Specifies to create new volumes.
- ◆ **Dynamic RDF Capability** — RDF capability of the volumes in the template. Possible values are:
 - **None** — Specifies to not use SRDF volumes.
 - **RDF Dynamic** — Specifies to only use dynamic SRDF volumes.
 - **R1** — Specifies to only use R1 volumes. This option is only available when you are creating/modifying a template that uses the Find Existing method for locating volumes; except when the volumes are thin (TDEVs), in which case this option is not available.
- ◆ **Emulation** — Emulation of the volumes in the template.

The following controls are available:

- ◆ **Create** — See [Creating storage templates](#) on page 252.
- ◆ **Delete** — See [Deleting storage templates](#) on page 255.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the storage group.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

External storage

Understanding external storage

External storage gives you the ability to use Federated Tiered Storage (FTS) to attach external storage to a Symmetrix system. Attaching external storage allows you to use physical disk space on existing Symmetrix systems while gaining access to Symmetrix features such as local replication, remote replication, storage tiering, data management, and data migration. In addition, FTS simplifies the management of federated multi-vendor or EMC storage systems.

For additional information about FTS, refer to the *EMC Symmetrix Federated Tiered Storage (FTS) Technical Notes* and the Federated Tiered Storage chapter of the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Virtualizing external LUNs

When you attach external storage to a Symmetrix system, FTS virtualizes an external Symmetrix system's SCSI logical units as Symmetrix disks called eDisks. eDisks have two modes of operation:

- ◆ **Encapsulation** — Allows you to preserve existing data on external Symmetrix systems and access it through Symmetrix volumes. These volumes are called encapsulated volumes.
- ◆ **External Provisioning** — Allows you to use external storage as raw capacity for new Symmetrix 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 is dependent on the external Symmetrix 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 Symmetrix system as disk group provisioned volumes.
- ◆ **Encapsulation for virtual provisioning (VP encapsulation)** — The eDisk is encapsulated and exported from the Symmetrix system as thin volumes.

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

External provisioning

When you virtualize an eDisk for external provisioning, you can then create Symmetrix 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.



If you use external provisioning, any data that is currently on the external volume is deleted.

Geometry of encapsulated volumes

Enginuity builds Symmetrix volumes based on the Symmetrix cylinder size (fifteen 64 K tracks), so the capacity of Symmetrix volumes will 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 Symmetrix cylinder size larger than the reported user-defined geometry are considered geometry limited. For additional details and a list of restrictions that apply to geometry-limited volumes, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

To virtualize external LUNs:

1. Select the Symmetrix system.
2. Select **Storage > External Storage** to open the **External Storage** page.
3. *Optional:* Click the **Not Virtualized** checkbox above the filtered list view to see a list of external LUNs that have not been virtualized.
4. Select the external LUNs that you want to virtualize.
5. Click **Virtualize** to open the **Virtualize External LUNs** dialog box.
6. Select an import method from the **Import Method** drop-down menu. This will determine the mode of operation for the eDisk.



If you select **Raw Space - External Provisioning**, any data that is currently on the external volume is deleted.

7. Select an external disk group from the **Disk Group** drop-down menu, or type a disk group name to create a new external disk group. Enginuity will add the virtualized external LUNs to the specified external disk group.
8. If you are using Virtual Provisioning, select an empty pool or an existing pool composed of externally provisioned data volumes from the **Thin Pool** drop-down menu. Type a pool name if you want to create a new pool.
9. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To override the auto meta member capacity configured on the Symmetrix system, specify the unit of measurement by selecting **GB**, **MB**, or **CYL** from the drop-down menu, and then select a capacity from the **Meta Member Capacity** drop-down menu.
The Total Enabled Pool Capacity in GB is displayed.
2. If you want all of the created Symmetrix volumes to be the same capacity, click the **Create Equal Meta Member Capacity** check box. If you do not select this check box, the meta tail will be smaller than the other volumes in the meta.
3. If you want to specify a DX director for the path to the eDisk, select a director from the **DX Director** drop-down menu.
4. Click **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.

Rescanning external storage

The first time you visit the External Storage page, Unisphere scans all of the volumes that are visible from the DX directors. When you return to the page, Unisphere populates the list views with cached data.

To rescan external storage:

1. Select the Symmetrix system.
2. Select **Storage > External Storage** to open the **External Storage** page.
3. Click **Rescan** to rescan the volumes that are visible from the DX directors. Note that the **Last External Storage Updated** field below the filtered list view updates with the data and time of the most recent scan.

Viewing External Storage

The External Storage page allows you to view and manage external storage as well as validate paths and zoning. The first time you visit the External Storage page, Unisphere scans all of the volumes that are visible from the DX directors.

EMC requires a minimum of four paths to external volumes, meaning that at least four ports belonging to a single DX dual initiator pair must be configured. The best practice for maximum redundancy is achieved by using single initiator/multiple target zoning. This is accomplished by creating individual zones that contain each DX port and all external ports that the external volumes are available on.

To view external storage and validate paths and zoning:

1. Select the Symmetrix system.
2. Select **Storage > External Storage** to open the **External Storage** page.

Use the tree view lists to filter the list of external LUNs by selecting various combinations of members within a tree list view (control ports, external ports, and external LUNs). You can select a single item, multiple items in consecutive rows (hold shift key and select), or multiple items in non-consecutive rows (hold control key and select). As each selection is made, the filtered results table is updated to reflect the current combination of filter criteria.

Control Ports tree view list

The following properties display:

- ◆ **Director** —Symmetrix system DX director.
- ◆ **Port** — Port number on the director.

External Ports tree view list:

The following properties display:

- ◆ **Port WWN** — World Wide Name of the external port.
- ◆ **Array ID** — External Symmetrix ID.
- ◆ **Vendor** — External Symmetrix vendor.

External LUNs tree view list:

The following properties display:

- ◆ **LUN WWN** — World Wide Name of the external LUN.
- ◆ **Capacity (GB)** — Capacity in GB of the external LUN.

Filtered results table:

The following properties display:

- ◆ **External LUN WWN** — World Wide Name of the external LUN.
- ◆ **Vendor** — Vendor name of the external LUN.
- ◆ **Capacity (GB)** — Capacity in GB of the external LUN.

- ◆ **Volume** —Volume ID on the external Symmetrix system.
- ◆ **LUN** — Displays 0 for EMC storage systems.
- ◆ **Virtualizing Status** — The mode of operation that the eDisk is using. Possible values are External, Encapsulated, and None.
- ◆ **Emulation** — Emulation type of the external LUN.
- ◆ **Disk Group** — Disk group that contains the virtualized LUN.
- ◆ **Spindle** — Spindle ID of the external spindle.
- ◆ **Service State** — Availability of the external LUN. Possible values are Normal, Degraded, and Failed. Failed means that there are no network paths available to the external LUN. Degraded means that there are paths from only one of the supporting DX directors. Normal means that there are network paths available from both supporting DX directors.

The following controls are available:

- ◆ **Virtualize** — See [Virtualizing external LUNs](#) on page 258.
- ◆ **Rescan** — See [Rescanning external storage](#) on page 260.

Reservations

Reserving volumes


The following explains how to reserve volumes and front-end mapping addresses for future configuration and masking operations. When using this feature, you reserve the volumes/addresses you plan on using, verify that no one else has reserved the resources, and release the reservations when the task is complete.

All reservations are assigned a reserve ID, indicating that the specified volumes/addresses are reserved. Any attempt to use the reserve volumes/addresses will return a message indicating that the volumes/addresses are reserved.

Before you begin:

Verify that the reservation feature is enabled for the Symmetrix system. For more information, see [See Setting system preferences](#) on page 32.

To reserve volumes:

1. Select the Symmetrix system.
2. Select **Storage > Reservations** to open the **Reservations** list view.
3. Click **Create** to open the **Create Reservation** dialog box.
4. Select the volumes.
5. Type any **Comments** regarding the reservation. The requirement for this field is set in system preferences.
6. Select **Expiration** and click calendar  to open the **Reservation Expiration** dialog box.
7. Select an **Expiration Date** and **Expiration Time**.
8. Click **OK** to close the **Reservation Expiration** dialog box.
9. Click **OK** to create the reservation.

Adding volumes to reservations

1. Select the Symmetrix system.
2. Select **Storage > Reservations** to open the **Reservations** list view.
3. Select the reservation and click **View Details** to open its **Details** view.
4. Click **Add** to open the **Add Volumes** dialog box.
5. Select one or more volumes and click **OK**.

Removing volumes from reservations

1. Select the Symmetrix system.
2. Select **Storage > Reservations** to open the **Reservations** list view.
3. Select the reservation and click **Release**.

4. Click **OK** in the confirmation message.

Releasing reservations

1. Select the Symmetrix system.
2. Select **Storage > Reservations** to open the **Reservations** list view.
3. Select one or more reservations and click **Release**.
4. Click **OK** in the confirmation message.

Viewing reservations

1. Select the Symmetrix system.
2. Select **Storage > Reservations** to open the **Reservations** list view.
Use the **Reservations** list view to display and manage reservations for a Symmetrix system.

The following properties display:

- ◆ **Reservation** — Reservation ID.
- ◆ **Owner** — User that created the reservation.
- ◆ **Application** — Application used to create the reservation.
- ◆ **Host** — Host from which the reservation was created.
- ◆ **Reserved Volumes** — Number of reserved volumes.
- ◆ **Creation** — Date/time the reservation was created.
- ◆ **Expiration** — Date/time the reservation will expire. Never is the default.
- ◆ **User Comment** — User-supplied comments.

The following controls are available:

- ◆ **Create** — [Reserving volumes](#) on previous page
- ◆ **View Detail** — See [Viewing reservation details](#) below.
- ◆ **Release** — [Releasing reservations](#) above

Viewing reservation details

1. Select the Symmetrix system.
2. Select **Storage > Reservations** to open the **Reservations** list view.
3. Select the reservation and click **View Details** to open its **Details** view.

Use the reservation **Details** view to display and manage a reservation. This view contains two panels, [Properties](#), and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Reservation** — Reservation ID.

- ◆ **Owner** — User that created the reservation.
- ◆ **Application** — Application used to create the reservation.
- ◆ **Host** — Host from which the reservation was created.
- ◆ **Reserved Volumes** — Number of reserved volumes.
- ◆ **Creation** — Date/time the reservation was created.
- ◆ **Expiration** — Date/time the reservation will expire. Never is the default.
- ◆ **User Comment** — User-supplied comments.

The following controls are available:

- ◆ **Create** — [Reserving volumes](#) on page 263
- ◆ **Release** — [Releasing reservations](#) on previous page
- ◆ **Add** — [Adding volumes to reservations](#) on page 263

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the reservation. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Reserved Volumes - 2** will open a view listing the two volumes held in the reservation.

Optimizer

Understanding Optimizer

Symmetrix Optimizer improves Symmetrix performance by continuously monitoring the back-end activity and swapping highly active volumes with idle volumes to achieve a balanced workload across the disks. This automated, continuous process is based on user-defined parameters and is completely transparent to end users, hosts, and applications in the environment.

For information on Optimizer concepts, refer to the *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*.

Managing Optimizer

The **Optimizer** dashboard provides you with a single place from which to manage Optimizer.

To access the Optimizer dashboard:

1. Select the Symmetrix system.
2. Select **Storage > Optimizer**.

Optimizer dashboard

The **Optimizer** dashboard provides access to the following Optimizer parameters:

◆ **Shared Parameters:**

The following parameters are shared between Optimizer and FAST. Therefore, changes made to these parameters will affect both Optimizer and FAST:

- **Data Movement Mode** — Specifies whether to move data automatically (Automatic) or to wait for user approval before moving data (User Approved).
- **Maximum Number of Volumes Moved Per Day** — Specifies the maximum number of moves to perform in a 24 hour period, starting at 12:00 AM. Possible values range from 2 to 200.
- **Maximum Simultaneous Volumes Moved** — Specifies the maximum number of moves that can be performed at one time. Possible values range from 2 to 32.

The following parameters are shared between Optimizer and FAST VP:

- **Workload Analysis Period** — Specifies the amount of workload sampling to maintain for sample analysis. Possible values are specified in units of time (hours, days, or weeks) and can range from 2 hours to 4 weeks, with the default being 8 hours.
- **Time to Sample before First Analysis** — Specifies the minimum amount of workload sampling to complete before analyzing the samples for the first time. When setting this parameter, be sure to allow enough time (usually a week) to establish a good characterization of the typical workload. This parameter allows you to begin operations before the entire Workload period has elapsed. Possible values range from 2 hours to the value specified for the Workload Analysis Period parameter.

◆ **Optimizer Specific Settings:**

- **Startup Mode** — Specifies whether Optimizer starts automatically with the Symmetrix service processor (Automatic), or if it requires manual intervention (Manual). This option does not apply/display for Symmetrix arrays running Enginuity 5874 or higher.
 - **Hot Spot Analysis** — Specifies whether to run the hot spot analysis algorithm. The hot spot analysis algorithm attempts to improve Symmetrix performance by swapping devices that it determines will reduce disk access times.
 - **Maintain Physical Mirror** — Specifies whether Optimizer should maintain any RAID 5 groups. This option does not apply/display for Symmetrix arrays running Enginuity 5874 or higher.
 - **Maintain RAID 5 Groups** — Specifies whether Optimizer should maintain any RAID 5 groups. This option does not apply/display for Symmetrix arrays running Enginuity 5874 or higher.
 - ◆ **DMX Lock Status:** Host that currently holds the Optimizer lock on the Symmetrix service processor. This option does not apply/display for Symmetrix arrays running Enginuity 5874 or higher.
 - ◆ **Optimizer Status:** Optimizer's current state.
- The **Optimizer** dashboard provides access to the following Optimizer controls:
- ◆ **Swap/Move List** — See [Viewing Optimizer swap/move lists](#) on page 269.
 - ◆ **Swap/Move History** — [Viewing Optimizer swap/move history](#) on page 270
 - ◆ **Approve Swap** — See [Approving Optimizer Swaps](#) on next page.
 - ◆ **Start/Lock** — [Lock/Unlock/Start/Stop dialog box](#) (Does not apply/display for Enginuity 5874 or higher).
 - ◆ **Enable/Disable** — [Enabling/Disabling Optimizer](#) on next page (Does not apply/display for versions lower than Enginuity 5874)
 - ◆ **Apply** — Applies changes made in the parameter list.
 - ◆ **Cancel** — Cancels changes made in the parameter list.

Starting/Stopping Optimizer

The following explains how to start/stop Optimizer on a Unisphere client.

Before you begin:

This procedure is only required when performing operations on a Symmetrix system running an Enginuity version lower than 5874.

To start/stop Optimizer:

1. Select the Symmetrix system.
2. Select **Storage** > **Optimizer** to open the **Optimizer** dashboard.
3. Click **Start/Lock** to open the **Lock/Unlock/Start/Stop** dialog box.
4. Click **Start** or **Stop**.

5. Click **OK**.
6. Click **OK** in the confirmation message.

Enabling/Disabling Optimizer

The following explains how to enable/disable Optimizer on a Symmetrix system running Enginuity 5874 or higher:

1. Select the Symmetrix system.
2. Select **Storage > Optimizer** to open the **Optimizer** dashboard.
3. Click **Enable/Disable** to open the **Enable/Disable** dialog box.
4. Click **Enable** or **Disable**.
5. Click **OK**.
6. Click **OK** in the confirmation message.

Locking/Unlocking Optimizer

The following explains how to lock/unlock Optimizer controls. Locking Optimizer prevents other users from making changes while you are working with the software.

Before you begin:

This procedure is only required when performing operations on a Symmetrix system running an Enginuity version lower than 5874.

To start/stop Optimizer:

1. Select the Symmetrix system.
2. Select **Storage > Optimizer** to open the **Optimizer** dashboard.
3. Click **Start/Lock** to open the **Lock/Unlock/Start/Stop** dialog box.
4. Click **Lock** or **Unlock**.
5. Click **OK**.
6. Click **OK** in the confirmation message.

Approving Optimizer Swaps

1. Select the Symmetrix system.
2. Select **Storage > Optimizer** to open the **Optimizer** dashboard.
3. Click **Approve Swap** to open the **Approve Swap** dialog box.

The following properties display, depending on the Enginuity version running on the Symmetrix system:

- ◆ **Group** — Group number associated with the swaps/moves.
- ◆ **Volume 1** — First volume involved in the swap/move.
- ◆ **Volume 2** — Second volume involved in the swap/move.

- ◆ **Status** — Status of the swap/move.
- ◆ **Execution Time** — Time the swap/move started.
- ◆ **Type** — Whether the plan was manually defined or auto generated by FAST or by Optimizer.

Engenuity versions lower than 5874:

- ◆ **Group** — Group associated with the swaps/moves.
- ◆ **Volume 1**— Volume being moved/swapped.
- ◆ **Mirror/Member** — Mirror position of the volume.
- ◆ **From** — Original location of the volume.
- ◆ **To** — New location of volume.
- ◆ **Status** — Status of the swap/move.
- ◆ **Execution Time** — Time the swap/move started.
- ◆ **Type** — Whether the plan was Manually Defined or Auto Generated by FAST or by Optimizer.

4. Select a swap.
5. Select the **Schedule for execution:** either **According to Optimizer policy** or specify a date and time.

Viewing Optimizer swap/move lists

The following explains how to view information on the *current* Optimizer/FAST plan:

1. Select the Symmetrix system.
2. Select **Storage > Optimizer** to open the **Optimizer** dashboard.
3. Click **Swap/Move List** to open the **Swap/Move List** dialog box.

The following properties display, depending on the Engenuity version running on the Symmetrix system:

Engenuity 5874 or higher:

- ◆ **Group** — Group associated with the swaps/moves.
- ◆ **Volume 1** — First volume involved in the swap/move.
- ◆ **Volume 2** — Second volume involved in the swap/move.
- ◆ **Status** — Status of the swap/move.
- ◆ **Execution Time** — Time the swap/move started.
- ◆ **Type** — Whether the plan was Manually Defined or Auto Generated by FAST or by Optimizer.
- ◆ **Group Percent Complete** — How much (%) of the plan is complete.

- ◆ **Estimated Time to Completion** — How long until the plan is complete.

Engenuity versions lower than 5874:

- ◆ **Group** — Group associated with the swaps/moves.
- ◆ **Volume** — Volume being moved/swapped.
- ◆ **Mirror/Member** — Mirror position of the volume.
- ◆ **From** — Original location of the volume.
- ◆ **To** — New location of volume.
- ◆ **Status** — Status of the swap/move.
- ◆ **Execution Time** — Time the swap/move started.
- ◆ **Type** — Whether the plan was Manually Defined or Auto Generated by FAST or by Optimizer.

Viewing Optimizer swap/move history

The following explains how to view information on the past Optimizer/FAST swaps/moves:

1. Select the Symmetrix system.
2. Select **Storage > Optimizer** to open the **Optimizer** dashboard.
3. Click **Swap/Move History** to open the **Swap/Move History** dialog box.
4. Specify the time range you are interested in and click **Run Filter**.

The following properties display, depending on the Engenuity version running on the Symmetrix system:

Engenuity 5874 or higher:

- ◆ **Group** — Group associated with the swaps/moves.
- ◆ **Volume 1** — First volume involved in the swap/move.
- ◆ **Volume 2** — Second volume involved in the swap/move.
- ◆ **Start Time**— Time (in Eastern Daylight Time) the swap/move started. EDT is 4 hours behind of Coordinated Universal Time (UTC).
- ◆ **End Time** — Time (in Eastern Daylight Time) the swap/move completed. EDT is 4 hours behind of Coordinated Universal Time (UTC).
- ◆ **Type** — Whether the plan was Manually Defined or Auto Generated by FAST or by Optimizer.

Engenuity versions lower than 5874:

- ◆ **Group** — Group associated with the swaps/moves.
- ◆ **Volume** — Volume being moved/swapped.
- ◆ **Mirror/Member** — Mirror position of the volume.

- ◆ **Hyper From** — Original location of the volume.
- ◆ **Hyper To** — New location of volume.
- ◆ **Attributes**— Whether the plan was Manually Defined or Auto Generated by FAST or by Optimizer.
- ◆ **Start Time**— Time (in Eastern Daylight Time) the swap/move started. EDT is 4 hours behind of Coordinated Universal Time (UTC).
- ◆ **End Time** — Time (in Eastern Daylight Time) the swap/move completed. EDT is 4 hours behind of Coordinated Universal Time (UTC).

CHAPTER 5

Host Management

This chapter contains the following:

- Initiators 274
- Host aliases 281
- Masking views 285
- Initiator groups 290
- Port groups 298
- Host Cache Adapters 303
- Virtual servers 304
- CU images 309

Initiators

Masking volumes

The following explains how to mask volumes on Symmetrix systems running Enginuity 5773 or lower. For instructions on masking volumes on Symmetrix systems running Enginuity 5874 or higher, refer to See [Creating masking views](#) on page 285..

You can perform this operation at the initiator or host alias level.

To mask volumes:

1. Select the Symmetrix system.
2. To mask at the initiator level:
 - a. Select **Hosts** > **Initiators** to open the **Initiators** list view.
 - b. Select an initiator, and click **Mask** to open the **Masking** dialog box.

To mask at the host alias level:

 - a. Select **Hosts** > **Host Aliases** to open the **Host Aliases** list view.
 - b. Select a host alias, and click **Mask** to open the **Masking** dialog box.
3. *Optional:* Select another **Director Port** for the operation. (One director at a time.)
4. *Optional:* Select another **Initiator/Alias** for the operation. (One director at a time.)
5. *Optional:* Select to **Include volumes not mapped to the port** in the **Available Volumes** list.
6. Select one or more **Available Volumes**, and click **Add** to move them to **Selected Volumes**, or click **Add All** to move all **Available Volumes** to the **Selected Volumes**.
7. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.
 - ◆ **Show Advanced** to set the advanced options, as described next.

Setting Advanced options:

1. To manually set the host LUN addresses:
 - a. Select the **Dynamic LUN Addressing** option.
 - b. Click **Set Dynamic LUN Address** to open the **Set Dynamic LUN Address** dialog box.
 - c. Select a volume, and notice the address displayed in the **Starting Host LUN** field. To accept this automatically generated address, click **Apply Starting LUN**. To move to the next available, click **Next Available LUN**.
 - d. Click **OK** once you are satisfied with the address.
2. To automatically map new volumes to all ports, select **Map Volumes**. (This feature is only available on Symmetrix DMX systems running Enginuity 5773.150 or higher.)

3. To refresh all the host-related profile data in the volume masking database (VCMDB), select **Refresh VCMDB after OK**.
4. Click **OK**.

Unmasking volumes

The following explains how to unmask volumes on Symmetrix systems running Enginuity 5773 or lower.

You can perform this operation at the initiator or host alias level.

1. Select the Symmetrix system.
2. To unmask at the initiator level:
 - a. Select **Host > Initiators** to open the **Initiator** list view.
 - b. Select the initiator from the list, and click **View Details** to open the **Details** view.
 - c. In the **Related Objects** panel, select **Volumes** to open the **Volumes** list view.

To unmask at the host alias level:

 - a. Select **Host > Host Aliases** to open the **Host Aliases** list view.
 - b. Select the alias and and click **View Details** to open its **Details** view.
 - c. In the **Related Objects** panel, select **Volumes** to open the **Volumes** list view.
3. Select a volume, and click **Unmask** to open the **Unmask Volumes** dialog box.
4. *Optional:* Click **Show Select Volumes** to view details on the selected volumes.
5. To unmap the volumes from their ports , select **Unmap Volumes**.
6. To refresh all the host-related profile data in the volume masking database (VCMDB), select **Refresh VCMDB after OK**.
7. Click **OK**.

Setting initiator port flags

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
3. Select an initiator, and click **Set Flags** to open the **Set Initiator Port Flags** dialog box.
4. *Optional:* Select an initiator whose flag settings you want to copy.
5. Modify the **attributes**, by selecting the corresponding **Override** option (thereby activating the **Enable** option), and enable (select) or disable (clear) the flag.
6. Click **OK**.

Setting initiator attributes

This procedure explains how to set an initiator's attributes.

Before you begin:

Any changes made to an initiator's attributes affect the initiator and all its ports.

To set initiator attributes:

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
3. Select an initiator, and click **Set Attributes** to open the **Set Initiator Attributes** dialog box.
The initiator director:port, initiator, and optional alias names display.
4. Type the **FCID (Fibre Channel ID) Value**.
5. To refresh all the host-related profile data in the volume masking database (VCMDB), select **Refresh VCMDB after OK**.
6. Click **OK**.

Renaming initiator aliases

When the system discovers the attached HBAs, a two-part record is created for the name. The format is NodeName/PortName. For fiber adapters, the HBA name is the WWN or iSCSI name. For native iSCSI adapters, the HBA name is the IP address.

You can rename the HBA identifier by creating a shorter, and easier to remember, ASCII alias name.

To rename an initiator alias:

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
3. Select an initiator, and click **Rename Alias** to open the **Rename Initiator Alias** dialog box.
4. Type a **Node Name** and **Port Name**.
5. Click **OK**.

This overwrites any existing alias name.

Replacing initiators

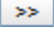
If a host adapter fails, or needs replacement for any reason, you can replace the adapter and assign its set of volumes to a new adapter.

To replace an initiator:

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
3. Select the initiator, and click **Replace Initiator** to open the **Replace Initiator Alias** dialog box.
The existing initiator and optional alias names display.
4. Type the full WWN or iSCSI identifier of the **New Initiator**. For native iSCSI, type the IP address.
5. Click **OK**.

This substitutes all occurrences of the old WWN/iSCSI/IP address with the new one.

Removing masking entries

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
There are two ways to delete initiator masking records, from the **Initiators view** or from the initiator **Details** view.
3. Select the initiator, click more , and select **Remove Masking Entry** to open the **Remove Masking Entry** dialog box.
4. Select the director and port .
5. To remove the entry from the Symmetrix login history tables, select **Include login history**.
6. To refresh all the host-related profile data in the volume masking database (VCMDB), select **Refresh VCMDB after OK**.
7. Click **OK**.

Viewing initiators

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
Use the **Initiators** list view to view and manage initiators.

The properties and controls displayed in the view vary depending on the Enginuity version running on the Symmetrix system and on how you arrived at this view.

- ◆ **Initiator** — WWN or IQN (iSCSI Qualified Name) ID of the initiator.
- ◆ **Dir:Port** — Symmetrix system director and port associated with the initiator, for example: FA-7E:1.
- ◆ **Alias** — User-defined initiator name.
- ◆ **Logged In** — Flag indicating if the initiator is logged into the fabric: Yes/No.
- ◆ **On Fabric** —Flag indicating if the initiator is on the fabric: Yes/No.
- ◆ **Port Flag Overrides** — Flag indicating if any port flags are overridden by the initiator: Yes/No.
- ◆ **Initiator Groups** — Number of initiator groups the initiator is associated with, including the immediate initiator group and any parent initiator groups that include this initiator group. This field only applies/appears for Symmetrix systems running Enginuity 5874 or higher.
- ◆ **Masking Views** — Number of masking views the initiator is associated with, including the masking views that are associated with any cascaded relationships. This field only applies/appears for Symmetrix systems running Enginuity 5874 or higher.
- ◆ **Port Flag Overrides** — Whether any of the initiator's port flags are overridden. This field only applies/appears for Symmetrix systems running Enginuity 5774 or lower.
- ◆ **Volumes** — Number of masked volumes. This field only applies/appears for Symmetrix systems running Enginuity 5774 or lower.

The following controls are available:

- ◆ **Add** — [Adding initiators to host aliases](#) on page 282
- ◆ **Remove** — [Removing initiators from host aliases](#) on page 282
- ◆ **Mask** — [Masking volumes](#) on page 281
- ◆ **Set Flags** — See [Setting initiator port flags](#) on page 275.
- ◆ **Set Attributes** — See [Setting initiator attributes](#) on page 275.
- ◆ **Rename Alias** — See [Renaming initiator aliases](#) on page 276.
- ◆ **Replace Initiator** — See [Replacing initiators](#) on page 276.
- ◆ **View Details** — See [Viewing initiator details](#) below.
- ◆ **Remove Masking Entry** — See [Removing masking entries](#) on previous page.

Viewing initiator details

1. Select the Symmetrix system.
2. Select **Hosts > Initiators** to open the **Initiators** list view.
3. Select the initiator from the list, and click **View Details** to open the initiator **Details** view.

Use the initiator **Details** view to view and manage initiators. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The properties and controls displayed in the view vary depending on the Enginuity version running on the Symmetrix system and on how you arrived at this view.

- ◆ **Initiator** — WWN or IQN (iSCSI Qualified Name) ID of the initiator.
- ◆ **Dir:Port** — Symmetrix system director and port associated with the initiator, for example: FA-7E:1
- ◆ **Alias** — The user-defined initiator name.
- ◆ **Number of Initiator Groups** — Number of associated initiator groups, including the immediate initiator group and any parent initiator groups that include this initiator group. This field only applies/appears for Symmetrix systems running Enginuity 5874 or higher.
- ◆ **Number of Masking Views** — Number of associated masking views, including the masking views that are associated with any cascaded relationships. This field only applies/appears for Symmetrix systems running Enginuity 5874 or higher.
- ◆ **Logged In** — Flag indicating if the initiator is logged into the fabric: Yes/No.
- ◆ **On Fabric** — Flag indicating if the initiator is on the fabric: Yes/No.
- ◆ **Port Flag Overrides** — Flag indicating if any port flags are overridden by the initiator: Yes/No.
- ◆ **Enabled Flags** — List of any enabled port flags overridden by the initiator.
- ◆ **Disabled Flags** — List of any disabled port flags overridden by the initiator.
- ◆ **Flags in Effect** — Flags that are in effect for the initiator.
- ◆ **Last Login** — Timestamp for the last time this initiator was logged into the system.
- ◆ **FCID** — Fibre Channel ID for the initiator.
- ◆ **FCID Value** — Value that is enabled for FCID lockdown.
- ◆ **FCID Lockdown** — Flag indicating if port lockdown is in effect: Yes/No.
- ◆ **IP Address** — IP address for the initiator.
- ◆ **LUN Offset** — Whether LUN offset is enabled. This feature allows you to skip over masked holes in an array of volumes.
- ◆ **Offset Value** — Number of address spaces required to skip over the hole.
- ◆ **Base Value** — Host's first missing LUN in the skip hole.
- ◆ **Dynamic Addressing** — Whether dynamic LUN addressing is enabled. With this feature, the system assigns the address based on availability.

The following controls are available:

- ◆ **Mask** — [Masking volumes](#) on page 281

- ◆ **Set Flags** — See [Setting initiator port flags](#) on page 275.
- ◆ **Set Attributes** — See [Setting initiator attributes](#) on page 275.
- ◆ **Rename Alias** — See [Renaming initiator aliases](#) on page 276.
- ◆ **Replace Initiator** — See [Replacing initiators](#) on page 276.
- ◆ **Removing Masking Entry** — See [Removing masking entries](#) on page 277.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the initiator. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Initiators Groups- 1** opens the view listing the initiator group that contains the initiator.

Host aliases

Masking volumes

The following explains how to mask volumes on Symmetrix systems running Enginuity 5773 or lower. For instructions on masking volumes on Symmetrix systems running Enginuity 5874 or higher, refer to See [Creating masking views](#) on page 285..

You can perform this operation at the initiator or host alias level.

To mask volumes:

1. Select the Symmetrix system.
2. To mask at the initiator level:
 - a. Select **Hosts>Initiators** to open the **Initiators** list view.
 - b. Select an initiator, and click **Mask** to open the **Masking** dialog box.

To mask at the host alias level:

 - a. Select **Hosts > Host Aliases** to open the **Host Aliases** list view.
 - b. Select a host alias, and click **Mask** to open the **Masking** dialog box.
3. *Optional:* Select another **Director Port** for the operation. (One director at a time.)
4. *Optional:* Select another **Initiator/Alias** for the operation. (One director at a time.)
5. *Optional:* Select to **Include volumes not mapped to the port** in the **Available Volumes** list.
6. Select one or more **Available Volumes**, and click **Add** to move them to **Selected Volumes**, or click **Add All** to move all **Available Volumes** to the **Selected Volumes**.
7. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.
 - ◆ **Show Advanced** to set the advanced options, as described next.

Setting Advanced options:

1. To manually set the host LUN addresses:
 - a. Select the **Dynamic LUN Addressing** option.
 - b. Click **Set Dynamic LUN Address** to open the **Set Dynamic LUN Address** dialog box.
 - c. Select a volume, and notice the address displayed in the **Starting Host LUN** field. To accept this automatically generated address, click **Apply Starting LUN**. To move to the next available, click **Next Available LUN**.
 - d. Click **OK** once you are satisfied with the address.
2. To automatically map new volumes to all ports, select **Map Volumes**. (This feature is only available on Symmetrix DMX systems running Enginuity 5773.150 or higher.)

3. To refresh all the host-related profile data in the volume masking database (VCMDB), select **Refresh VCMDB after OK**.
4. Click **OK**.

Creating host aliases

The following explains how to create host aliases:

1. Select the Symmetrix system.
2. Select **Hosts > Host Aliases** to open the **Host Aliases** list view.
3. Click **Create** to open the **Create Host Alias** dialog box.
4. Type a **Host Alias Name** (up to 16 alpha-numeric characters).
5. Select one or more **Initiators** and click **OK**.

Adding initiators to host aliases

1. Select the Symmetrix system.
2. Select **Hosts > Host Aliases** to open the **Host Aliases** list view.
3. Select the host alias, and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Initiators** to open the **Initiators** list view.
5. Click **Add** to open the **Add Initiator to Host Alias** dialog box.
6. Select one or more initiators and click **OK**.

Removing initiators from host aliases

1. Select the Symmetrix system.
2. Select **Hosts > Host Aliases** to open the **Host Aliases** list view.
3. Select the host alias, and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Initiators** to open the **Initiators** list view.
5. Select one or more initiators and click **Remove**.
6. Click **OK** in the confirmation message.

Unmasking volumes from host aliases

1. Select the Symmetrix system.
2. Select **Hosts > Host Aliases** to open the **Host Aliases** list view.
3. Select the host alias, and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Volumes** to open the **Volumes** list view.
5. Select one or more volumes and click **Unmask**.
6. To automatically unmap the volumes from all ports, select **Unmap Volumes**.
7. To refresh all the host-related profile data in the volume masking database (VCMDB), select **Refresh VCMDB**.
8. Click **OK** in the confirmation message.

Deleting host aliases

The following explains how to delete host aliases:

1. Select the Symmetrix system.
2. Select **Hosts** > **Host Aliases** to open the **Host Aliases** list view.
3. Select one or more aliases and click **Delete**.
4. Click **OK** in the confirmation message.

Renaming host aliases

When a Symmetrix system discovers the attached HBAs, a two-part record is created of the name. The format is HostName/HBAname. For fiber adapters, the HBA name is the WWN or iSCSI name. For native iSCSI adapters, the HBA name is the IP address.

You can rename the HBA identifier by creating a shorter, and easier to remember alias name:

1. Select the Symmetrix system.
2. Select **Hosts** > **Host Aliases** to open the **Host Aliases** list view.
3. Select an alias and click **Rename Alias** to open the **Rename Alias** dialog box.
4. Type the two-part Initiator Alias (HostName/NewAdapterAlias). For example: **api172/20000000920dc290** to **api172/p2**.
5. Click **OK**.

This operation will overwrite any existing alias name.

Viewing host aliases

1. Select the Symmetrix system.
 2. Select **Hosts** > **Host Aliases** to open the **Host Aliases** list view.
- Use the **Host Aliases** list view to display and manage host aliases.

The following properties display

- ◆ **Host Alias Name** — The name of the host alias.
- ◆ **Number of Initiators** — The number of masking records in the host alias.
- ◆ **Number of Volumes** — The number of volumes masked to the initiators.

The following controls are available:

- ◆ **Mask** — See [Masking volumes](#) on page 281.
- ◆ **Create** — See [Creating host aliases](#) on previous page.
- ◆ **Set Flags** — See [Setting initiator port flags](#) on page 275.
- ◆ **View Details** — See [Viewing host alias details](#) on next page.
- ◆ **Delete Host Alias** — See [Deleting host aliases](#) above.

Viewing host alias details

1. Select the Symmetrix system.
2. Select **Hosts** > **Host Aliases** to open the **Host Aliases** list view.
3. Select the host alias, and click **View Details** to open its **Details** view.

Use the host alias **Details** view to display and manage a host alias. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Host Alias Name** — the name of the host alias.
- ◆ **Number of Initiators** — The number of masking records in the host alias..
- ◆ **Number of Volumes** — The number of volumes masked to the initiators .

The following controls are available:

- ◆ **Mask** — See [Masking volumes](#) on page 281.
- ◆ **Set Flags** — See [Setting initiator port flags](#) on page 275.
- ◆ **Delete Host Alias** — See [Deleting host aliases](#) on previous page.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the host alias. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Initiators - 2** opens a view listing the two initiators in the host alias.

Masking views

Creating masking views

The following explains how to mask volumes on Symmetrix systems running Enginuity 5874 or higher. For instructions on masking volumes on Symmetrix systems running Enginuity 5773 or lower, refer to See [Masking volumes](#) on page 281..

Before you begin:

- ◆ To create a masking view you need to have created initiator groups, port groups, and storage groups.
For instructions, refer to [Creating initiator groups \(Hosts\)](#) on page 290, [Creating port groups](#) on page 298, and [Creating storage groups](#) on page 148.
- ◆ You can also use the **Host Management - Create Host** wizard available from **Common Tasks > Create a new host** to create groups.

To create a masking view:

1. Select the Symmetrix system.
2. Select **Hosts > Masking view** to open the **Masking view** list view.
There are two ways to create an masking view, from the **Masking view** list view or from the **Details** view.
3. Click **Create** to open the **Create Masking View** dialog box.
4. Type the **Masking View Name**. Masking view names must be unique from other masking views on the array and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Masking view names are case-insensitive.
5. Select the **Initiator Group**.
6. Select the **Port Group**.
7. Select the **Storage Group**.
8. Click **OK**.

Renaming masking views

1. Select the Symmetrix system.
2. Select **Hosts > Masking view** to open the **Masking Views** list view.
3. Select the masking view from the list and click **View Details** to open the **Details** view.
4. Type the new **Name**, and click **Apply**. Masking view names must be unique from other masking views on the array and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Masking view names are case-insensitive.

Deleting masking views

1. Select the Symmetrix system.
2. Select **Hosts > Masking view** to open the **Masking View** list view.
There are two ways to delete a masking view, from the **Masking view** list view or from the Masking view **Details** view.
3. Select the masking view from the list, click **Delete** to open the **Delete Masking View** confirmation dialog box.
4. To unmap volumes in the masking view from their mapped ports, select **Unmap**.
5. Click **OK**.

Viewing masking views

1. Select the Symmetrix system.
2. Select **Hosts > Masking view** to open the **Masking Views** list view.
Use the **Masking view** list view to view and manage masking views.
The following properties display:
 - ◆ **Name** — User-defined masking view name.
 - ◆ **Initiator Group** — Name of the associated initiator group.
 - ◆ **Port Group** — Name of the associated port group.
 - ◆ **Storage Group** — Name of the associated storage group.The following controls are available:
 - ◆ **Create** — See [Creating masking views](#) on previous page.
 - ◆ **View Details** — See [Viewing masking view details](#) on page 288.
 - ◆ **View Connections** — See [Viewing masking view connections](#) below.
 - ◆ **Delete** — See [Deleting masking views](#) above.

Viewing masking view connections

This procedure explains how to view masking view connections.

To view masking view connections:

1. Select the Symmetrix system.
2. Select **Hosts > Masking view** to open the **Masking view** list view.
3. Select the masking view from the list and click **View Connections** to open the masking view connections view.

Use the **Masking view** view to filter a masking view by selecting various combinations of members within a group (initiators, ports, volumes) and display the masking view details from the group level to the object level.

Filtering a masking view

The **Masking view** view contains three tree view lists for each of the component groups in the masking view, initiator groups, ports groups, and storage groups.

The parent group is the default top-level group in each expandable tree view and contains a list of all components in the masking group including child entries which are also expandable.

To filter the masking view, single or multi-select (hold shift key and select) the items in the list view. As each selection is made the filtered results table is updated to reflect the current combination of filter criteria.

Masking view tree view lists

The following properties display:

- ◆ **Initiator Group** — Name of top-level initiator group.
- ◆ **Initiator/Alias** — WWN or IQN (iSCSI Qualified Name) ID and the alias of the initiator group.
- ◆ **Port Group** — Name of top-level port group.
- ◆ **Director:Port** — Symmetrix system director and port in the port group.
- ◆ **Storage Group** — Name of top-level storage group.
- ◆ **Volume** — Symmetrix system volume number.
- ◆ **Capacity (GB)** — Capacity, in GB, of the volume.

Filtered results table

The following properties display:

- ◆ **LUN Address** — LUN address number.
- ◆ **Volume** — Symmetrix system volume number.
- ◆ **Capacity (GB)** — Capacity, in GB, of the volume.
- ◆ **Initiator/Alias** — WWN or IQN (iSCSI Qualified Name) ID and alias of the initiator.
- ◆ **Dir:Port** — Symmetrix system director and port in the port group.
- ◆ **Logged In** — Indicates if the initiator is logged into the host/target.
- ◆ **On Fabric** — Indicates if the initiator is zoned in and on the fabric.

The following additional filters are available to filter the results table:

- ◆ **Show logged in** — Shows only the entries for LUNs where the associated initiator is logged in.
- ◆ **Show On Fabric** — Shows only the entries for LUNs where the associated initiator is zoned in and on the fabric.

Viewing masking view details

1. Select the Symmetrix system.
2. Select **Hosts > Masking view** to open the **Masking Views** list view.
3. Select the masking view from the list and click **View Details** to open the **Details** view.
Use the **Details** view to view and manage masking views. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Name** — User-defined masking view name.
- ◆ **Initiator Group** — Name of the associated initiator group.
- ◆ **Number of initiators** — Number of initiators in the masking view.
This is the number of primary initiators contained in the masking view and does not include any initiators included in cascaded initiator groups that may be part of the masking view.
- ◆ **Port Group** — Name of the associated port group.
- ◆ **Number of ports** — Number of ports contained in the masking view.
- ◆ **Storage Group** — Name of the associated storage group.
- ◆ **Number of volumes** — Number of volumes in the storage group contained in the masking view.
- ◆ **Capacity (GB)** — Total capacity, in GB, of all volumes in the masking view.

The following controls are available:

- ◆ **Create** — See [Creating masking views](#) on page 285.
- ◆ **Delete** — See [Deleting masking views](#) on page 286.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the masking view.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the masking view. Each group link is followed the name of the group, or by a number, indicating the number of objects in the corresponding view. For example, clicking **Volumes - 2** opens the view listing the two volumes associated with the masking view.

Initiator groups

Creating initiator groups (Hosts)

This procedure explains how to create initiator groups (hosts) using the **Host Management - Create Host** wizard. There are two ways to open the wizard, from an initiator's **Details** view (described next), or from the **Create a new host** option in the **Common Tasks** panel.

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. *Optional:* Select an initiator group and click **View Details** to open its **Details** view.
4. Click **Create** to open the **Create Host** wizard.
5. Select the **Symmetrix** system on which to create the group.
6. Type a name for the **Host** (name for the initiator group). Initiator group names must be unique from other initiator groups on the array and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Initiator group names are case-insensitive.
7. Select an initiator or initiator group to associate with the host and click **Add**. Note that you can filter the list by starting to type the name. In addition, you can also use an asterisk (*) as a wildcard. Repeat this step for each initiator/group you want associate with the host.

Note Initiators can only belong to one initiator group at a time; therefore, any initiators that do not appear in the list are already part of a group.

8. Click one of the following:
 - ◆ **Finish** to create the initiator group.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To override the port attributes, do the following:
 - a. Click **Set Initiator Group Flags** to open the **Set Initiator Group Flags** dialog box.
 - b. *Optional:* Select an initiator group whose flag settings you want to copy.
 - c. *Optional:* Modify any of the [attributes](#), by selecting the corresponding **Override** option (thereby activating the **Enable** option) and enable (select) or disable (clear) the flag.
 - d. *Optional:* Select **Consistent LUNs** to specify that LUN values for the initiator group must be kept consistent for all volumes within each masking view of which this initiator group is part. When set, any masking operation involving this initiator group that would result in inconsistent LUN values, will be rejected. When not set, the Symmetrix system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations.
 - e. Click **OK** to close the **Set Initiator Group Flags** dialog box.
2. To create a cascaded initiator group, select **Auto-Cascade**.

When this option is selected Unisphere for VMAX will create an initiator group for each individual initiator, and then add these new groups to the parent initiator group.

Naming conventions for the **Auto-Cascade** feature are described in the following example:

If the **Host Name** entered is "server1" and there are 2 HBAs/Initiators selected, the parent initiator group is named "server1", the first HBA is placed in an initiator group named server1_1 and the second HBA is placed in an initiator group named server1_2. The result is a parent initiator group "server1" with two child initiator groups "server1_1" and server1_2".

3. Click **Finish**.

Adding/Removing initiators from initiator groups

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Initiators** to open the **Initiator** list view.
5. Do the following, depending on whether you are adding or removing initiators:

Adding initiators:

- a. Click **Add** to open the **Add Initiators to Initiator Group** dialog box.
- b. Select the initiators and click **Add**.
- c. Click **OK**.

Removing initiators:

- a. Select one or more initiators and click **Remove**.
- b. Click **OK** in the confirmation message.

Adding/Removing initiator groups from initiator groups

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Child IGs** to open the **Child IGs** list view.
5. Do the following, depending on whether you are adding or removing initiator groups:

Adding initiator groups:

- a. Click **Add** to open the **Add Initiator Groups** dialog box.
- b. Select an initiator group and click **Add**. Repeat this step for each additional initiator group.
- c. Click **OK**.

Removing initiator groups:

- a. Select one or more initiator groups and click **Remove**.
- b. Click **OK** in the confirmation message.

Modifying initiator groups

There are two methods for modifying an initiator group, [from the Initiator Groups list view](#), or from the **Host Management** wizard, located in the **Common Tasks** panel.

From the Initiator Group view:

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group and click **View Details** to open its **Details** view.
4. To change the name of the initiator group, highlight the name in the **Properties** panel, and type a new name over it.
5. To override the port attributes, do the following:
 - a. Click **Set Flags** to open the **Set Initiator Group Flags** dialog box.
 - b. *Optional:* Select an initiator group whose flag settings you want to copy.
 - c. *Optional:* Modify any of the **attributes**, by selecting the corresponding **Override** option (thereby activating the **Enable** option) and enable (select) or disable (clear) the flag.
 - d. *Optional:* Select **Consistent LUNs** to specify that LUN values for the initiator group must be kept consistent for all volumes within each masking view of which this initiator group is part. When set, any masking operation involving this initiator group that would result in inconsistent LUN values, will be rejected. When not set, the Symmetrix system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations.
 - e. Click **OK** to close the **Set Initiator Group Flags** dialog box.
6. Click **Apply**.

From the Manage Host wizard:

1. In the **Common Task** panel, click **Manage host** to open the **Host Management** wizard.
2. Select the **Symmetrix** system on which the group is defined.
3. Select the **Host** (name for the initiator group).
4. Click **Next**.
5. To add an initiator or initiator group, select one and click **Add**.
To remove an initiator or initiator group, select one and click **Remove**.
6. Click one of the following:
 - ◆ **Finish** to create the initiator group.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To override the port attributes, do the following:
 - a. Click **Set Initiator Group Flags** to open the **Set Initiator Group Flags** dialog box.
 - b. *Optional:* Select an initiator group whose flag settings you want to copy.
 - c. *Optional:* Modify any of the **attributes**, by selecting the corresponding **Override** option (thereby activating the **Enable** option) and enable (select) or disable (clear) the flag.

- d. *Optional:* Select **Consistent LUNs** to specify that LUN values for the initiator group must be kept consistent for all volumes within each masking view of which this initiator group is part. When set, any masking operation involving this initiator group that would result in inconsistent LUN values, will be rejected. When not set, the Symmetrix system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations.
- e. Click **OK** to close the **Set Initiator Group Flags** dialog box.
- f. Click **Finish**.

Deleting initiator groups

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
There are two ways to delete an initiator group, from the **Initiator Groups** list view or from the **Details** view.
3. Select the initiator group from the list, click **Delete** to open the **Delete Initiator Group** confirmation dialog box.
4. Click **OK**.

Renaming initiator groups

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group from the list and click **View Details** to open the initiator group **Details** view.
4. Type the initiator group **Name**, and click **Apply**. Initiator group names must be unique from other initiator groups on the array and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and hyphens (-) are allowed. Initiator group names are case-insensitive.

Setting initiator group flags

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group and click **Set Flags** to open the **Set Initiator Group Flags** dialog box.
4. *Optional:* Select an initiator group whose flag settings you want to copy.
5. Modify any of the **attributes**, by selecting the corresponding **Override** option (thereby activating the **Enable** option) and enable (select) or disable (clear) the flag.
6. *Optional:* Select **Consistent LUNs** to specify that LUN values for the initiator group must be kept consistent for all volumes within each masking view of which this initiator group is part. When set, any masking operation involving this initiator group that would result in inconsistent LUN values, will be rejected. When not set, the Symmetrix system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations.
7. Click **OK**.







Viewing initiator groups

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.

Use **Initiator Groups** list view to view and manage initiator groups.

There are multiple ways to open this view. Depending on the one you used, some of the following properties and controls may not appear.

The following properties display:

- ◆ **Name** — User-defined initiator group name, accompanied by an icon indicating the initiator group type. Possible types are:
 -  Standalone initiator group
 -  Child initiator group
 -  Parent initiator group
 -  Standalone initiator group in masking view
 -  Child initiator group in masking view
 -  Parent initiator group in masking view
- ◆ **Parent IGs** — Number of initiator groups where the initiator is a member.
- ◆ **Child IGs** — Number of child initiator groups in the initiator group.
- ◆ **Masking Views** — Number of masking views where the initiator group is associated.
- ◆ **Initiators** — Number of initiators in the group including initiators in any child initiator groups.
- ◆ **Consistent LUNs** — Flag indicating if the Consistent LUNs flag is set.
- ◆ **Port Flag Overrides** — Flag indicating if any port flags are overridden by the initiator: Yes/No.
- ◆ **Last Update** — Timestamp of the most recent changes to the initiator group.

The following controls are available:

- ◆ **Create** — See [Creating initiator groups \(Hosts\)](#) on page 290.
- ◆ **Set Flags** — See [Setting initiator group flags](#) on previous page.
- ◆ **View Details** — See [Viewing initiator group details](#) below.
- ◆ **Delete** — See [Deleting initiator groups](#) on previous page.

Viewing initiator group details

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group from the list and click **View Details** to open the **Details** view.

Use the initiator group **Details** view to view and manage the initiator group. This view contains two panels, the [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Name** — User-defined initiator group name.
- ◆ **Child IGs** — Number of child initiator groups in the initiator group.
- ◆ **Parent IGs** — Number of parent initiator groups where the initiator group is a member.
- ◆ **Initiators** — Number of initiators in the group including initiators in any child initiator groups.
- ◆ **Masking Views** — Number of masking views where the initiator group is associated.
- ◆ **Port Flag Overrides** — Flag indicating if any port flags are overridden by the initiator group: Yes/No.
- ◆ **Enabled Port Flags** — List of any enabled port flags overridden by the host.
- ◆ **Disabled Port Flags** — List of any disabled port flags overridden by the host.
- ◆ **Last Update** — Timestamp of the most recent changes to the initiator group.

The following controls are available:

- ◆ **Create** — See [Creating initiator groups \(Hosts\)](#) on page 290.
- ◆ **Set flags** — See [Setting initiator group flags](#) on page 293.
- ◆ **Delete** — See [Deleting initiator groups](#) on page 293.
- ◆ **Apply** — Applies new initiator group name entered in the **Name** field.
- ◆ **Cancel** — Cancels the rename action.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the initiator group. Each group link is followed the name of the group, or by a number, indicating the number of objects in the corresponding view. For example, clicking **Initiators - 2** opens the view listing the two initiators contained in the initiator group.

Viewing initiators in initiator group

1. Select the Symmetrix system.
2. Select **Hosts > Initiator Groups** to open the **Initiator Groups** list view.
3. Select the initiator group from the list and click **View Details** to open the initiator group **Details** view.
4. From the **Related Objects** panel, click **Initiators** to open the initiator group **Initiators** list view.

The following properties display:

- ◆ **Name** — WWN or IQN (iSCSI Qualified Name) ID of the initiator.
- ◆ **Alias** — User-defined initiator name.
- ◆ **Masking Views** — Number of associated masking views.

The following controls are available:

- ◆ **Add** — See [Adding/Removing initiators from initiator groups](#) on page 291.
- ◆ **Remove** — See [Adding/Removing initiator groups from initiator groups](#) on page 291.
- ◆ **View Details** — See [Viewing initiator details](#) on page 278.

Port attributes

Table 6: Port attributes

Attribute	Description
Common Serial Number	Enables multi-path configurations or hosts that need a unique serial number to determine which paths lead to the same volume.
Volume Set Addressing	Enables the volume set addressing mode. When using volume set addressing, you must specify a 4-digit address in the following range: (0)000-(0)007, (0)010-(0)017,... to a maximum of (0)FF0-(0)FF7 Where the first digit must always be set to 0 (Symmetrix does not currently support the upper range of volume set addressing), the second digit is the VBus number, the third digit is the target, and the fourth digit is the LUN. This field does not appear for Symmetrix systems running Enginuity 5874 or higher.
Avoid Reset Broadcast	Enables a SCSI bus reset to only occur to the port that received the reset (not broadcast to all channels).
Environ Set	Enables the environmental error reporting by the Symmetrix to the host on the specific port.

Attribute	Description
Disable Q Reset on UA	When enabled, a Unit Attention (UA) that is propagated from another director does not flush the queue for this volume on this director. Used for hosts that do not expect the queue to be flushed on a 0629 sense (only on a hard reset).
SCSI 3	Alters the inquiry data (when returned by any volume on the port) to report that the Symmetrix supports the SCSI-3 protocol. When disabled, the SCSI 2 protocol is supported.
SCSI Support1 (OS2007)	Provides a stricter compliance with SCSI standards for managing volume identifiers, multi-port targets, unit attention reports, and the absence of a volume at LUN 0. To enable the SCSI Support1 attribute, you must also enable the SPC2 Protocol Version attribute.
SPC2 Protocol Version	This flag should be enabled (default) in a Windows 2003 environment running Microsoft HCT test version 12.1. When setting this flag, the port must be offline.
Sunapee	Enables the Sunapee option on the port, for SUN PDB clusters. This field does not appear for Symmetrix systems running Enginuity 5874 or higher.
Siemens	Returns in the sense data error 0B48 instead of 0B44 for normal behavior. For Siemens R-series platforms only. This field does not appear for Symmetrix systems running Enginuity 5874 or higher.
Sequent	Sets the task timeout window to be 15 seconds before aborting a process, and the Symmetrix system on the specified port does not force wide or synchronous negotiations. Also, a busy status returns, instead of a 0B44h when aborting a command on a timeout. This field does not appear for Symmetrix systems running Enginuity 5874 or higher.
Enable AS400	Indicates whether AS/400 is enabled.
Open VMS	Enables an Open VMS fiber connection.

Port groups

Creating port groups

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
3. Click **Create** to open **Create Port Group** dialog box.
4. Type a **Port group name**. Port group names must be unique from other port groups on the Symmetrix system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Port group names are case-insensitive.
5. Select the available ports from **Ports** list, and click **Add** to add them to the **Ports to add** list.
The following properties display:
 - ◆ **Dir:Port** — Symmetrix system director and port in the port group.
 - ◆ **Ports** — Number of port groups where the port is a member.
 - ◆ **Masking Views** — Number of masking views where the port is associated.
 - ◆ **Volumes** — Number of volumes in the port group.
 - ◆ **VSA Flag** — An indicator to show if Volume Set Addressing is set for the port.
6. Click **OK**.

Deleting port groups

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
There are two ways to delete a Port Group, from the **Port Groups** list view or from the port group **Details** view.
3. Select the port group and click **Delete**.to open the **Delete Port Group** confirmation message.
4. *For mapped ports only:* Select **Unmap**.
5. Click **OK**.

Adding ports to port group

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
3. Select the port group and click **View Details** to open the port group **Details** view.
4. From the **Related Objects** panel, click **Ports** to open the **Ports** in port group list view.
5. Click **Add** to open the **Add Ports** dialog box.
6. Select the available ports from the **Ports to add** list, and click **Add** to add them to the **Ports to Add** list.

The following properties display:

- ◆ **Dir:Port** — Symmetrix system director and port in the port group.
 - ◆ **Ports** — Number of port groups where the port is a member.
 - ◆ **Masking Views** — Number of associated masking views.
7. Click **OK**.

Removing ports from port group

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
3. Select the port group and click **View Details** to open the port group **Details** view.
4. From the **Related Objects** panel, click **Ports** to open the **Ports from Ports Group** list view.
5. Select the port to remove or hold down the shift key to multi-select the ports to be removed from the port group.
6. Click **Remove** to open the **Remove Ports** confirmation message.
7. *For mapped ports only:* Select **Unmap**.
8. Click **OK**.

Renaming port groups

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
3. Select the port group and click **View Details** to open the initiator group **Details** view.
4. Type the new port group **Name** and click **Apply**.

Viewing port groups


1. Select the Symmetrix system.
2. Select **Hosts** > **Port Groups** to open the **Port Groups** list view.


The **Port Groups** list view allows you to view and manage port groups on a Symmetrix system.

There are multiple ways to open this view. Depending on the one you used, some of the following properties and controls may not appear.

The following properties display (Click a column heading to sort the list by that value):

- ◆ **Name** — User-defined port group name, accompanied by an icon indicating the port group type. Possible types are:

 Standalone port group

 Port group in masking view

- ◆ **Ports** — Number of ports in the group.
- ◆ **Masking Views** — Number of masking views where the port group is associated.
- ◆ **Last Update** — Timestamp of the most recent changes to the port group.

The following controls are available:

- ◆ **Create** — See [Creating port groups](#) on page 298.
- ◆ **View Details** — See [Viewing port groups details](#) below.
- ◆ **Delete** — See [Deleting port groups](#) on page 298.

Viewing port groups details

1. Select the Symmetrix system.
2. Select **Hosts** > **Port Groups** to open the **Port Groups** list view.
3. Select the port group and click **View Details** to open the port group **Details** view.
4. Use the port groups **Details** view to view and manage a port group. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display (Click a column heading to sort the list by that value):

- ◆ **Name** — User-defined port group name.
- ◆ **Number of Ports** — Number of ports in the group.
- ◆ **Number of Masking Views** — Number of masking views where the port group is associated.
- ◆ **Last Update** — Timestamp of the most recent changes to the port group.
- ◆ **Host I/O Limit (I/O/Sec)** — Total host I/O limit on the specified port group in IO/Sec. Zero indicates that there is no limit set.

- ◆ **Host I/O (MB /Sec)** — Total host I/O limit on the specified port group in MB/Sec. Zero indicates that there is no limit set.
- ◆ **Negotiated Speed (MB/Sec)** — Bandwidth in MB/sec for that port group (that is, the aggregated port negotiated speed for the ports in the group).
- ◆ **Percent Capacity** — Percentage of the bandwidth demand over the port group negotiated speed.
- ◆ **Excess (MB/Sec)** — Amount of bandwidth in MB/sec that is left available on the port group after the host I/O limits have been accounted for.

The following controls are available:

- ◆ **Create** — See [Creating port groups](#) on page 298.
- ◆ **Delete** — See [Deleting port groups](#) on page 298.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the port group.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the port group. Each group link is followed the name of the group, or by a number, indicating the number of objects in the corresponding view. For example, clicking **Host I/O Limits - 2** opens the view listing the host I/O limits set for the two storage groups associated with the port group.

Viewing ports in port group

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
3. Select the port group and click **View Details** to open its **Details** view.
4. From the **Related Objects** panel, click **Ports** to open the **Ports** list view.

Use the **Ports** list view to view and manage ports.

The following properties are displayed:

- ◆ **Dir:Port** — Symmetrix system director and port in the port group.
- ◆ **Port Groups** — Number of port groups where the port is a member.
- ◆ **Masking Views** — Number of masking views where the port is associated.
- ◆ **Mapped Volumes** — Number of volumes mapped to the port.

The following controls are available:

- ◆ **Add** — See [Adding ports to port group](#) on page 299.
- ◆ **Remove** — See [Removing ports from port group](#) on page 299.
- ◆ **View Details** — See [Viewing port details](#) on next page.

Viewing port details

1. Select the Symmetrix system.
2. Select **Hosts > Port Groups** to open the **Port Groups** list view.
3. Select the port group and click **View Details** to open its **Details** view.
4. From the **Related Objects** panel, click **Ports** to open the **Ports** list view.
5. Select a port and click **View Details** to open its **Details** view.

Use the port **Details** view to view and manage a port. This view contains two panel, Properties and Related Objects.

Properties panel

The following properties display:

- ◆ **Dir:Port** — Symmetrix system director and port in the port group.
- ◆ **Number of Port Groups** — Number of port groups where the port is a member.
- ◆ **Number of Masking Views** — Number of masking views where the port is associated.
- ◆ **Number of Masked Volumes** — Number volumes visible through the port.
- ◆ **Number of Mapped Volumes** Number of volumes mapped to the port, including meta members.
- ◆ **Volume Set Addressing** — Whether volume set addressing is on or off.
- ◆ **Port Status** — Whether the port is online or offline.

Related Objects panel

The **Related Objects** panel links you to views displaying objects associated with the port. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Masking Views - 2** will open a view listing the two masking views associated with the port.

Host Cache Adapters

Viewing host cache adapters

1. Select the Symmetrix system.
2. Select **Hosts** > **Host Cache Adapters** to open the **Host Cache Adapters** list view.

The following properties display:

- ◆ **Card S/N** — Adapter serial number.
- ◆ **Card Version** — Adapter version.
- ◆ **Vender** — Adapter vender.
- ◆ **Card Size (GB)** — Adapter size.
- ◆ **Card Used (GB)** — Amount of card used.
- ◆ **Volumes** — Number of accessible volumes.
- ◆ **Host** — Host name.
- ◆ **IP Address** — Host IP address.
- ◆ **Host OS** — Host operating system.

Virtual servers

Adding a new virtual server

This procedure explains how to register a virtual server (VMware or Hyper V). You must complete this procedure before you can add the virtual server's storage to the VM.

Before you begin:

- ◆ To perform this operation, you must be an Administrator or StorageAdmin.
- ◆ For more information on integrating Unisphere for VMAX in a virtual environment, refer to the Techbook, *Using EMC Symmetrix Storage in VMware vSphere Environments*.

To add a new virtual server:

1. Select the Symmetrix system or **All Symmetrix** systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.
3. Click **Add** to open the **Add New Server** dialog box.
4. Type the fully qualified name or IP address of the server (**Server/IP Address**).
5. Type the **User Name** used to connect to the server. The user must have root privileges on the server.
6. Type the **Password** used to connect to the server.
7. Retype the password for confirmation.
8. Select the type of server to add (**VMware** or **Hyper-V**).
9. *Optional:* Select **Retrieve Info** to have Unisphere retrieve all the information on the server. If Unisphere returns an error message when retrieving information on an ESX server, verify that the ESX server's domain name matches the name in the credential file.
10. Click **OK**.

Adding storage to a VM

This procedure explains how to add a virtual server's storage to a VM.

Before you begin:

To perform this operation, you must be an Administrator or StorageAdmin.

To add a new virtual server:

1. Select the Symmetrix system or **All Symmetrix** systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.
3. Select the virtual server and click **View Details** to open its **Details** view.
4. Click **Volumes - *nn*** in the **Related Objects** panel to open the **Volumes** list view.
5. Select one or more volumes and click **Add VM Storage** to open the **Add VM Storage** dialog box.
6. Select the **VM** to which you are adding the volumes.
7. Click **OK**.

Removing a virtual server

This procedure explains how to remove a virtual server.

Before you begin:

To perform this operation, you must be an Administrator or StorageAdmin.

To add a new virtual server:

1. Select the Symmetrix system or All Symmetrix systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.
3. Select the virtual server and click **Remove**.
4. Click **OK** in the confirmation message.

Note A virtual server can also be removed from the virtual server Details view.

Removing storage from a VM

This procedure explains how to remove storage from a virtual server.

Before you begin:

To perform this operation, you must be an Administrator or StorageAdmin.

To add a new virtual server:

1. Select the Symmetrix system or All Symmetrix systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.
3. Select the virtual server and click **View Details** to open the **Details** view.
4. Click **Volumes - nn** in the **Related Objects** panel to open the **Volumes** list view.
5. Select one or more volumes and click **Remove VM Storage**.
6. Click **OK** in the confirmation message.

Changing the password on a virtual server

This procedure explains how to change the password used to access a virtual server.

Before you begin:

To perform this operation, you must be an Administrator or StorageAdmin.

To add a new virtual server:

1. Select the Symmetrix system or **All Symmetrix** systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.
3. Select the virtual server and click **Change Password**.
4. Click **OK** on the confirmation message to open the **Change Virtual Server Password** dialog box.
5. Type a new **Password** and repeat it for confirmation.
6. Click **OK**.

Note The virtual server password can also be changed from the virtual server **Details** view.

Viewing virtual servers

1. Select the Symmetrix system or **All Symmetrix** systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.

The following properties display:

- ◆ **Server/IP Address** — The fully-qualified server name or IP Address.
- ◆ **Server Type** — The virtual server type. Possible values are VMware and Hyper-V.
- ◆ **VMs** — The number of VMs.
- ◆ **Volumes** — The number of volumes.
- ◆ **Last Updated** — The timestamp of the refresh for the virtual server. If the server does not have a timestamp, click **View Details** to rescan the server.

The following controls are available:

- ◆ **Add** — See [Adding a new virtual server](#) on page 304.
- ◆ **View Details** — See [Viewing the details of a virtual server](#) below.
- ◆ **Change Password** — See [Changing the password on a virtual server](#) on previous page.
- ◆ **Remove** — See [Removing a virtual server](#) on page 305.

Viewing the details of a virtual server

1. Select the Symmetrix system or All Symmetrix systems.
2. Select **Hosts > Virtual Servers** to open the **Virtual Servers** list view.
3. Select the virtual server and click **View Details** to open its **Details** view.

This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Server/IP Address** — The fully-qualified server name or IP Address.
- ◆ **Server Type** — The virtual server type. Possible values are VMware and Hyper-V.
- ◆ **Total Memory** — Total memory of the virtual server.
- ◆ **Build** — The virtual server's build number.
- ◆ **Version** — The virtual server's version number.
- ◆ **Last Updated** — The timestamp of the last refresh of the virtual server.

The following controls are available:

- ◆ **Change Password** — See [Changing the password on a virtual server](#) on previous page.
- ◆ **Remove** — See [Removing a virtual server](#) on page 305.

Related Objects panel

The Related Objects panel provides links to views for the objects contained in or associated with the virtual server. Each link provides the name of the related object and the number of items. For example, clicking **Volumes - 34** opens the **Volumes** list view showing the 34 volumes related to the virtual server.

CU images

Mapping CKD volumes

The following explains how to map CKD volumes to ESCON/FICON ports.

You can perform this operation at the volume level or the CU image level.

1. Select the Symmetrix system.
2. To map at the volume level:
 - a. Select **Storage > Volumes** to open the **Volume Dashboard**.
 - b. In the **Volumes** panel, select the type of CKD volume.

Hint To display only CKD volumes in the **Volumes** panel, set the **Emulation** filter to **CKD**.

- c. Click **View** to open the **CKD Volumes** list view.
- d. Select a CKD volume, and click **z/OS Map** to open the **z/OS Map Volumes** dialog box.

To map at the CU image level:

- a. Select **Hosts > CU Images** to open the **CU Images** list view.
- b. Select an image and click **z/OS Map** to open the **z/OS Map** dialog box.
3. Type or **Select a Volume Range**.
4. Type or **Select the Base Address** to be assigned to the first volume in the mapping request. It increases incrementally by one for each volume in the range of volumes being mapped.
5. Type or **Select an SSID**. Valid SSIDs must only have unmapped volumes using them and the number of volumes cannot exceed 256.
6. Select the **Port** to which you want to map the volumes.
7. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Unmapping CKD volumes

The following explains how to unmap CKD volumes from ESCON/FICON ports.

You can perform this operation at the volume level or the CU image level.

1. Select the Symmetrix system.
2. To unmap at the volume level:
 - a. Select **Storage > Volumes** to open the **Volume Dashboard**.
 - b. In the **Volumes** panel, select the type of CKD volume.

Hint To display only CKD volumes in the **Volumes** panel, set the **Emulation** filter to **CKD**.

- c. Click **View** to open the **CKD Volumes** list view.

- d. Select a CKD volume, and click **z/OS Unmap** to open the **z/OS Unmap Volumes** dialog box.

To unmap at the CU image level:

- a. Select **Hosts > CU Images** to open the **CU Images** list view.
 - b. Select an image and click **z/OS Unmap** to open the **Unmap CU Image** dialog box.
3. Type or **Select** the **Volume Range** to be unmapped.
 4. Type or **Select** the **Base Address**.
 5. Type or **Select** an SSID. Valid SSIDs must only have unmapped volumes using them, and the number of volumes cannot exceed 256.
 6. Select the **Port** to which you want to map the volumes.
 7. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.


Assigning alias addresses

The following explains how to assign PAV aliases to mapped CKD volumes on a Symmetrix system running Enginuity 56xx or lower. For instructions on assigning aliases on Symmetrix systems running 5771 or higher, refer to [Assigning an alias range](#).

Before you begin:


PAV alias must be enabled on the Symmetrix system.

To assign an alias address:

1. Select a Symmetrix system.
2. Select **Hosts > CU Images** to open the **CU Images** list view.
3. Select an image, click more  , and select **Assign Alias Addresses** to open the **Assign Alias Addresses** dialog box.
4. Type or **Select** a **Volume Range**.
5. Type a **Starting Alias** address.
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Removing alias addresses

The following explains how to remove alias addresses from CKD volumes on a Symmetrix system running Enginuity 56xx or earlier. For instructions on removing aliases on Symmetrix systems running 5771 or higher, refer to See [Assigning alias ranges](#) below.

1. Select a Symmetrix system.
2. Select **Hosts > CU Images** to open the **CU Images** list view.
3. Select an image, click more , and select **Remove Alias Addresses** to open the **Remove Alias Addresses** dialog box.
4. Type or **Select a Volume Range**. Alternately, you can also select a different CU image.
5. Type a **Starting Alias**.
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.


Assigning alias ranges

The following explains how to assign a range of PAV aliases to mapped CKD volumes on a Symmetrix system running 5771 or higher. For instructions on assigning aliases on Symmetrix systems running Enginuity 56xx or lower, refer to See [Assigning alias addresses](#) on previous page.

Before you begin:

PAV alias must be enabled on the Symmetrix system.

To assign an alias address:

1. Select a Symmetrix system.
2. Select **Hosts > CU Images** to open the **CU Images** list view.
3. Select an image, click more , and select **Assign Alias Range** to open the **Assign Alias Range** dialog box.
4. Type the **Starting Alias** address.
5. Type the **Ending Alias** address.
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.


Removing alias ranges

The following explains how to remove aliases from CKD volumes on a Symmetrix system running 5771 or higher. For instructions on removing aliases on Symmetrix systems running Engenuity 56xx or earlier, refer to See [Removing alias addresses](#) on previous page.

Before you begin:


PAV alias must be enabled on the Symmetrix system.

To remove alias ranges:

1. Select a Symmetrix system.
2. Select **Hosts** > **CU Images** to open the **CU Images** list view.
3. Select an image, click more , and select **Remove Alias Range** to open the **Remove Alias Range** dialog box.
4. Type the **Starting Alias** address.
5. Type the **Ending Alias** address.
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.


Assigning alias counts

The following explains how to add a number of aliases to mapped CKD volumes on Symmetrix systems running Engenuity 57xx or higher:

1. Select a Symmetrix system.
2. Select **Hosts** > **CU Images** to open the **CU Images** list view.
3. Select an image, click more , and select **Add Alias Count** to open the **Add Alias Count** dialog box.
4. Type or **Select** the **Volume Range** to which the aliases will be added.
5. Type the number of aliases to add to each volume in the range (**Alias Count**).
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Removing alias counts

The following explains how to remove a number of aliases from mapped CKD volumes on Symmetrix systems running Engenuity 57xx or higher:

1. Select a Symmetrix system.
2. Select **Hosts > CU Images** to open the **CU Images** list view.
3. Select an image, click more  , and select **Remove Alias Count** to open the **Remove Alias Count** dialog box.
4. Type or **Select** the **Volume Range** from which to remove the aliases.
5. Type the number of aliases to remove from each volume in the range (**Alias Count**).
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing CU images

1. Select the Symmetrix system.
2. Select **Hosts > CU Images** to open the **CU Images** list view.

Use the **CU Images** list view to view and manage CU images.

The following properties display:

- ◆ **Image SSID** — CU image number 0x0-0xFF.
- ◆ **SSID** — SSID (subsystem ID) for the image.
- ◆ **Number of Ports** — The number of ports to which the CU image is mapped.
- ◆ **Number of Volumes** — The number of volumes in the CU.
- ◆ **Number of Base Addresses** — The number of base addresses assigned.
- ◆ **Number of Aliases** — The number of aliases assigned.
- ◆ **Status** — The status of volumes in the CU image.

The following controls are available:

- ◆ **View Details** — See [Viewing CU image details](#) on next page.
- ◆ **z/OS Map** — See [Mapping CKD volumes](#) on page 309.
- ◆ **z/OS Unmap** — See [Unmapping CKD volumes](#) on page 309.
- ◆ **Assign Alias Addresses** — See [Assigning alias addresses](#) on page 310.
- ◆ **Remove Alias Addresses** — See [Removing alias addresses](#) on page 311.
- ◆ **Assign Alias Range** — See [Assigning alias ranges](#) on page 311.
- ◆ **Remove Alias Range** — See [Removing alias ranges](#) on previous page.
- ◆ **Assign Alias Count** — See [Assigning alias counts](#) on previous page.
- ◆ **Remove Alias Count** — See [Removing alias counts](#) on previous page.

Viewing CU image details

1. Select the Symmetrix system.
2. Select **Hosts** > **CU Images** to open the **CU Images** list view.
3. Select the CU image, and click **View Details** to open its **Details** view.

Use the CU image **Details** view to display and manage a CU image. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Image SSID** — CU image number 0x0-0xFF.
- ◆ **SSID** —SSID (subsystem ID) for the image.
- ◆ **Number of Ports** — The number of ports to which the CU image is mapped.
- ◆ **Number of Volumes** — The number of volumes in the CU.
- ◆ **Number of Base Addresses** — Number of base addresses assigned.
- ◆ **Number of Aliases** — The number of aliases assigned.
- ◆ **Status** — The status of volumes in the CU image.

The following controls are available:

- ◆ **z/OS Unmap** — See [Unmapping CKD volumes](#) on page 309.
- ◆ **Assign Alias Addresses** — See [Assigning alias addresses](#) on page 310.
- ◆ **Remove Alias Addresses** — See [Removing alias addresses](#) on page 311.
- ◆ **Assign Alias Range** — See [Assigning alias ranges](#) on page 311.
- ◆ **Remove Alias Range** — See [Removing alias ranges](#) on page 312.
- ◆ **Assign Alias Count** — See [Assigning alias counts](#) on page 312.
- ◆ **Remove Alias Count** — See [Removing alias counts](#) on page 312.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the CU image. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Volumes - 2** opens a view listing the two volumes in the image.

CHAPTER 6

Data Protection

This chapter contains the following:

- Local Replication 316
- Device Groups 351
- Remote Replication 357
- Migration 384
- Replication Groups and Pools 389
- RecoverPoint 409
- Open Replicator 415
- Federated Live Migration 425

Local Replication

Monitoring local replication sessions

The **Local Replication** dashboard provides you with a single place to monitor and manage local replication sessions on a Symmetrix system.

To access the Local Replication dashboard:

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication**.

Local Replication dashboard

The **Local Replication** dashboard consists of three views, one for each of the local replication types: **TimeFinder/Clone** and **TimeFinder/Snap** and **TimeFinder/Mirror**.

TimeFinder/Clone view

The **TimeFinder/Clone** view allows you to view and manage groups containing volumes using TimeFinder Clone.

The following properties display:

- ◆ **Source Group** — Lists the groups containing volumes using TimeFinder Clone. Information in this column is organized in a tree format, with groups organized into folders according to their type. To view information on a specific group, expand the appropriate folder.
- ◆ **Standard** — The number of standard volumes in the group.
- ◆ **BCV** — The number of BCVs in the group.
- ◆ **Target** — The number of target volumes in the group.
- ◆ **State** — The combined state of the sessions in the group. If all the sessions are in the same state, then that state appears; otherwise, **Mixed** appears.
- ◆ **Group Type** — The type of group. Possible values are RDF1, RDF2, RDF21, and Regular.
- ◆ **Group Valid** — Indicates whether the group is valid. The possible values are Yes or No.

The following controls are available:

- ◆ **Create Sessions** —See [Creating clone copy sessions](#) on page 319.
- ◆ **Activate** —See [Activating clone copy sessions](#) on page 321.
- ◆ **Recreate** —See [Recreating clone copy sessions](#) on page 321.
- ◆ **View Details** —See [Viewing clone copy sessions](#) on page 329.
- ◆ **Set Mode** —See [Modifying clone copy sessions](#) on page 324.
- ◆ **Terminate** —See [Terminating clone copy sessions](#) on page 328.
- ◆ **Establish** —See [Establishing clone copy sessions](#) on page 322.
- ◆ **Restore** —See [Restoring data from target volumes](#) on page 326.
- ◆ **Split** —See [Splitting clone volume pairs](#) on page 327.

TimeFinder/Snap view

The **TimeFinder/Snap** view allows you to view and manage groups containing volumes using TimeFinder/Snap.

The following properties display:

- ◆ **Source Group** — Lists the groups containing volumes using TimeFinder/Snap. Information in this column is organized in a tree format, with groups organized into folders according to their type. To view information on a specific group, expand the appropriate folder.
- ◆ **Standard** — The number of standard volumes in the group.
- ◆ **BCV** — The number of BCVs in the group.
- ◆ **VDEV** — The number of virtual volumes in the group.
- ◆ **Target** — The number of target volumes in the group.
- ◆ **State** — The combined state of the sessions in the group. If all the sessions are in the same state, then that state appears; otherwise, **Mixed** appears.
- ◆ **Group Type** — The type of group. The possible values are RDF1, RDF2, RDF21, and Regular.
- ◆ **Group Valid** — Indicates whether the group is valid. Possible values are Yes or No.

The following controls are available:

- ◆ **Create** — See [Creating virtual copy sessions](#) on page 333.
- ◆ **Activate** — See [Activating virtual copy sessions](#) on page 334.
- ◆ **Terminate** — See [Terminating virtual copy sessions](#) on page 340.
- ◆ **View Details** — See [Viewing snap pair details](#) on page 341.
- ◆ **Duplicate** — See [Duplicating virtual copy sessions](#) on page 337.
- ◆ **Establish** — See [Establishing virtual copy sessions](#) on page 335.
- ◆ **Recreate** — See [Recreating virtual copy sessions](#) on page 338.
- ◆ **Restore** — See [Restoring virtual copy sessions](#) on page 338.

TimeFinder/Mirror

The **TimeFinder Mirror** view allows you to view and manage groups containing volumes using TimeFinder/Mirror.

The following properties display:

- ◆ **Source Group** — Lists the groups containing volumes using TimeFinder/Mirror. Information in this column is organized in a tree format, with groups organized into folders according to their type. To view information on a specific group, expand the appropriate folder.
- ◆ **Standard** — The number of standard volumes in the group.
- ◆ **BCV** — The number of BCVs in the group.
- ◆ **State** — The combined state of the sessions in the group. If all the sessions are in the same state, then that state appears; otherwise, **Mixed** appears.
- ◆ **Group Type** — The type of group. Possible values are RDF1, RDF2, RDF21, and Regular.
- ◆ **Group Valid** — Indicates whether the group is valid. The possible values are Yes or No.

The following controls are available:

- ◆ **Establish** — See [TimeFinder/Mirror](#) on page 343.
- ◆ **Restore** — See [Restoring BCV pairs](#) on page 344.
- ◆ **Split** — See [Splitting BCV pairs](#) on page 345.
- ◆ **View Details** — [Viewing mirror pair details](#)
- ◆ **Cancel** — See [Cancelling BCV pairs](#) on page 346.

TimeFinder/Clone

Understanding TimeFinder/Clone operations

Clone copy sessions allow 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, as long as 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.

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 session's progress.](#)
5. [Terminate the session.](#)

For more information on TimeFinder/Clone concepts, refer to the *EMC Solutions Enabler Symmetrix TimeFinder Family CLI Product Guide* and the *EMC TimeFinder Family Product Guide* available at: support.emc.com.

Creating clone copy sessions

This procedure explains how to create clone copy sessions.

Before you begin:

- ◆ You can only perform this operation on a group containing source and target volumes.
- ◆ You can use the target volume of a clone session as the source volume for other clone sessions. To use this feature, you must first enable the SYMAPI_ALLOW_DEV_IN_MULT_GRP option in the SYMAPI options file. For more information on enabling SYMAPI options, refer to the *Solutions Enabler Symmetrix CLI Command Reference* available at: support.emc.com.
- ◆ The clone copy does not become available to the host until the session is activated.

To create clone copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.

3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

Select a group and click **Create Sessions** to open the **Create** dialog box.

Pair level:

 - a. Select a group, and click **View Details** to open its **Details** view.
 - b. Select one or more pairs, and click **Create Sessions** to open the **Create Sessions** dialog box.
 - c. Click **Set Pairs** to open the **Set Pairs** dialog box.
 - d. Select a source volume, and a target volume and click **Add** to make them a pair. Repeat this step as required.
 - e. Click **OK** to return to the **Create Sessions** dialog box.
5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To select a **Pairing Type**, expand the menu and select one of the following. If you are not using the **Pairing Type** option, leave this field set to **None**.
 - ◆ **Use Exact Pairs** — Allows the system to pair up the volumes in the exact order that they were added to the group.
 - ◆ **Use Optimized Pairs** — Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.
2. To select a **Copy Mode**, expand the menu, and select one of the following:
 - ◆ **Use VP SNAP** — Specifies to create the session with TimeFinder VP Snap, which allows multiple sessions to share allocations within a thin pool, thus reducing the storage required for saved tracks.
 - ◆ **Use No Copy** — Specifies to change the session to CopyOnAccess once the session is activated and no full-volume copy will initiate.
 - ◆ **Use Pre-Copy** — Specifies to start copying tracks in the background before you activate the clone session.
3. By default, when creating a clone session, the system will create an SDDF session for maintaining changed track information. To change this default behavior, expand the **Differential Mode** menu, and select **Use No Differential**. Otherwise, leave this field set to **Use Differential**.

4. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
5. Click **OK**.

Activating clone copy sessions

This procedure explains how to activate the copy operation from the source volume to the target volume. Activating the copy session places the target volume in the Read/Write state. The target host can access the cloned data and has access to data on the source host until you terminate the copy session.

Before you begin:

You can only activate clone sessions that are in the Created or Recreated state.

To activate clone copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

Select a group and click **Activate** to open the **Activate** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, and click **Activate** to open the **Activate Sessions** dialog box.

5. Select the type of **Source Volumes** to use.
6. Select the type of **Target Volumes** to use.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Recreating clone copy sessions

This procedure explains how to incrementally copy all subsequent changes made to the source volume (made after the point-in-time copy initiated) to the target volume.

While in the Recreated state, the target volume remains Not Ready to the host.

Before you begin:

- ◆ You can recreate a session once a clone volume is in the CopyInProgress state to incrementally copy all subsequent changes made to the source volume.
- ◆ The copy session must not have been created with the No Copy or No Differential option.
- ◆ The session must have been activated to establish the new point-in-time copy.
- ◆ With Engenuity 5876 Q4 2012 SR and higher, you can recreate a clone copy without terminating TimeFinder/Snap or VP Snap sessions that are cascading off of the clone target.


To recreate clone copy sessions:


1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Recreate** to open the **Recreate** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more , and select **Recreate** to open the **Recreate Sessions** dialog box.

5. Select the clone session, click more , and select **Recreate** to open the **Recreate** dialog box.
6. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
7. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
8. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Establishing clone copy sessions

This procedure explains how to create and immediately activate clone copy sessions.

Before you begin:

The establish operation sets the target volume to Not Ready for a short time. If you are using a file system, unmount the target host before performing the establish operation.


To establish clone copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Establish** to open the **Establish** dialog box.

Pair level:

- a. Select a group and click **View Details** to open its **Details** view.
 - b. Select one or more pairs, click more , and select **Establish** to open the **Establish Sessions** dialog box.
5. Specify whether to perform an **Incremental** or **Full** establish.
 6. If performing this operation at the pair level, do the following:
 - a. Click **Set Pairs** to open the **Set Pairs** dialog box.
 - b. Select a source volume and a target volume and click **Add** to make them a pair. Repeat this step as required.
 - c. Click **OK** to return to the **Establish Sessions** dialog box.
 7. Select the **Source Type**.
 8. Select the **Target Type**.
 9. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To specify **Full Establish Options**, expand the menu, and do the following:
 - a. Select one of the following options. If you are not using the **Full Establish Options**, leave this field set to **None**.
 - ◆ **Use Exact Pairs** — Allows the system to pair up the volumes in the exact order that they were added to the group.
 - ◆ **Use Optimized Pairs** — Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.
 - b. To create the session with TimeFinder VP Snap, select **Use VP Snap**. Using this option allows multiple sessions to share allocations within a thin pool, thus reducing the storage required for saved tracks.

2. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
3. Click **OK**.

Modifying clone copy sessions

This procedure explains how to modify the mode in which a clone copy session is operating.

Before you begin:

- ◆ You can modify the mode between Copy, NoCopy, and Precopy on clone pairs that are in a Created, Recreated, or Activated state.
- ◆ Do not attempt to change a session created with the Differential option to the No Copy mode, as the session will fail.

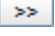
To modify clone copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Set Mode** to open the **Set Mode** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more , and select **Set Mode** to open the **Set Mode** dialog box.

5. Select a **Copy Mode**:

- ◆ **Copy** — If the session was created without the Copy option, it can be changed now to Copy mode. A copy initiates once the session is activated.
- ◆ **No Copy** — If the session was created with Copy mode, you can change the session to NoCopy mode. The session becomes CopyOnAccess once the session is activated and no full-volume copy will initiate.
- ◆ **Precopy** — If the session was created without Precopy, you can change the session to Precopy mode, which implies a copy. You cannot change to NoCopy mode. Once the session is activated, the session changes to Copy mode.

6. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
7. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
8. Click one of the following:
 - ◆ **OK** to start the operation now.

- ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Restoring data from target volumes

This procedure explains how to copy target data to another volume (full restore), or back to the original source volume (incremental restore).

In the case of a full restore, the original session terminates and a copy session to the target of the restore starts.

In the case of an incremental restore, the original session copy direction is reversed and changed data is copied from the target volume to the source volume. To support this operation, the session must have been created with the Differential option and the volume must be in a fully Copied state.

Before you begin:

- ◆ For a clone session in the Created state, the target volume must be in a fully copied state.
- ◆ With Enginuity 5875 or higher, you can:
 - Use ORS control volumes as clone restore targets when the volumes are in PUSH sessions and in the ORS Copied state.
 - Perform an incremental restore to a cascaded clone target. For example, in the relationship A->B->C, you can copy data from volume C to volume A.
- ◆ With Enginuity 5875 Q2 2011 SR, you can perform an incremental restore on volume pairs in a NoCopy/NoDiff clone session.
- ◆ With Enginuity 5876 Q4 2012 SR, you can perform an incremental restore of clone targets to source volumes with active snap and VP snap sessions.

To restore data from a target volume:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

Select a group, and click **Restore** to open the **Restore** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, and click **Restore** to open the **Restore Sessions** dialog box.

5. Select a **Restore Type**:
 - ◆ **Full Restore** — Terminates the original session and starts a copy session to the target of the restore.
 - ◆ **Incremental Restore** — Terminates the original session and starts an incremental copy session back to the original source volume. The session must have been created with the Differential option.
6. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. Select a **Clone Copy Mode**.
2. Select how to create the clones (**Clone Creation**). To allocate pool capacity for a region of the thin volumes, type a starting and ending cylinder number.
3. Click either of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Splitting clone volume pairs

This procedure explains how to split clone volume pairs. Splitting volume pairs changes the direction of the clone relationship (that is, the original source volume becomes the source volume for a future copy), which enables you to use either the establish or recreate command.

Before you begin:


- ◆ The clone session must be in the Restored state.

To split clone volume pairs:

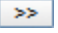
1. Select the Symmetrix system.
2. Select **Data Protection** › **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.

4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Split** to open the **Split** dialog box.

Pair level:

- a. Select a group and click **View Details** to open its **Details** view.
 - b. Select one or more pairs, click more , and select **Split** to open the **Split Sessions** dialog box.
5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
 6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
 7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Terminating clone copy sessions

This procedure explains how to terminate a clone copy session, thereby deleting the pairing information from the Symmetrix system, and removing any hold on the target volume.

Before you begin:

- ◆ You need a clone copy session in any pair state.
- ◆ Terminating a session while the pairs are in the CopyOnAccess, CopyOnWrite, or CopyInProgress state causes the session to end. If the application has not finished accessing all of the data, the target copy is not a full copy.

To terminate clone copy sessions:

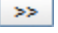
1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.

4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Terminate** to open the **Terminate** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more , and select **Terminate Sessions** to open the **Terminate Sessions** dialog box.

5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Viewing clone copy sessions

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Clone** to open the **TimeFinder Clone** view.
4. Select a group, and click **View Details** to open its **Clone Pair List** view.

Use the **Clone Pair List** view to display and manage a group's clone pairs.

The following properties display:

- ◆ **Group Details** — Lists additional group details, including:
 - **Group Type** — The type of group. Possible values are: Regular, R1, R2, or R21.
 - **Vender ID** — The company who's application created the group.
 - **Group Valid** — Whether the group is valid (**Yes**) or invalid (**No**).
 - **Group Creation Time** — Date/Time the group was created.
 - **Application ID** — The application that created the group.
 - **Group Modify Time** — Date/Time the group was last modified.

- ◆ **Session List** — Lists the group's clone pairs and their attributes, including:
 - **Source Volume** — The name of the source volume.
 - **Source LDev** — The logical name of the source volume.
 - **Target Volume** — The name of the target volume.
 - **Target LDev** — The logical name of the target volume.
 - **State** — The session state of the pair.
 - **CDGP** — Flags specific to the pair session in the form:
 - (C): X = The background copy setting is active for this pair.
 . = The background copy setting is not active for this pair.
 - (G): X = The Target volume is associated with a group.
 . = The Target volume is not associated with a group.
 - (D): X = The Clone session is a differential copy session.
 . = The Clone session is not a differential copy session.
 - (P): X = The precopy operation has completed one cycle.
 . = The precopy operation has not completed one cycle.
 - **Percent Copied** — The percentage of copying that is complete.

The following controls are available:

- ◆ **Create Sessions** — See [Creating clone copy sessions](#) on page 319.
- ◆ **Activate** — See [Activating clone copy sessions](#) on page 321.
- ◆ **Recreate** — See [Recreating clone copy sessions](#) on page 321.
- ◆ **View Details** — [Viewing clone session details](#)
- ◆ **Set Mode** — See [Modifying clone copy sessions](#) on page 324.
- ◆ **Terminate** — See [Terminating clone copy sessions](#) on page 328.
- ◆ **Establish** — See [Establishing clone copy sessions](#) on page 322.
- ◆ **Restore** — See [Restoring data from target volumes](#) on page 326.
- ◆ **Split** — See [Splitting clone volume pairs](#) on page 327.

Clone copy session options

The following table describes the TimeFinder/Clone session options:

Table 7: TimeFinder/Clone session options

Session option	Description	Available with action
Both Sides	Activates all locally and remotely associated clone pairs in an SRDF group.	Activate Establish

Session option	Description	Available with action
Concurrent	Performs the action for an additional clone pair in a group.	Create Recreate Establish Activate Verify
Consistent	Creates clone copies that are consistent with the database up to the point in time that the activation occurs. It suspends writes to the source volumes during the activation.	Activate
Copy	Creates a full data copy. By omitting this option (default), the volume pair state will be in the CopyOnAccess state when activated. Actual copying of the data is deferred until either tracks on the source volume are written to, or tracks on the target volume are read or written. This option is only applicable when the target volume is a regular volume (not a virtual volume).	Create Establish
Differential	Used with either the Copy or Precopy option to create an SDDF session for maintaining changed track information. It must be used when creating copy sessions on which you plan on issuing a Restore action.	Create Establish
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.	Create Establish Activate Restore Split Terminate
Not Ready	Sets the target volumes as Not Ready.	Establish Activate Restore
Optimize	Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups. For remote volumes, use the Optimize Rag option.	Create Establish
Optimize Rag	Uses optimization rules to create remote BCV pairs from volumes within the same RDF (RA) group on a Symmetrix system.	Create Establish
Precopy	Copies tracks in the background before the clone session is	Create

Session option	Description	Available with action
	activated. Used with the create and recreate actions.	Recreate
Restored	With the verify command, verifies that the copy sessions are in the Restored state. With the terminate command, terminates a restored VP Snap session.	Verify Terminate
Star	Targets the action at volumes in SRDF/Star mode.	Create Recreate Establish Activate Restore Split Terminate
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.	Terminate

TimeFinder/Snap

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 Symmetrix 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 associated with the session is freed and returned to the SAVE volume pool for future use.

The following are the basic actions 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 new 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, refer to the *EMC Solutions Enabler Symmetrix TimeFinder Family CLI Product Guide* and the *EMC TimeFinder Family Product Guide* at [EMC Online Support Site](#).

Creating virtual copy sessions

Virtual copy sessions define and set up the volumes for snap operations.

The Create action defines the copy session requirements and sets the track protection bitmap on the source volume to protect all tracks and detect which tracks are being accessed by the target host or written to by the source host. The target virtual volume remains Not Ready to its host and placed on hold status for copy session usage. This prevents other control operations from using the volume. The volume pair state transitions from CreateInProg to Created when complete. The virtual data becomes accessible to its host when the copy session is activated.

Before you begin:

- ◆ You can create up to 128 copies of a source volume to various virtual target volumes. This requires that you enable the following SYMCLI environment variable:
SYMCLI_MULTI_VIRTUAL_SNAP = ENABLED
- ◆ A source volume can concurrently copy data to as many as 15 target volumes at one time. Each target requires a separate copy session.
- ◆ For Symmetrix systems running Enginuity 5875 or higher, you can:
 - Use this feature to create multivirtual snap sessions from thin volumes.
 - Use RDF2 async volumes as source volumes.
 - Create a snap pair from a clone target in the Split state.
- ◆ To create a snap session of an R2 volume that is in an SRDF/A session, volume level pacing must be enabled on the R1 side.

To create virtual copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.

4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

 - a. Select a group, and click **Create Sessions** to open the **Create** dialog box.
 - b. Select a **Snap Pool**.

Pair level:

 - a. Select a group, and click **View Details** to open its **Details** view.
 - b. Select one or more pairs, and click **Create Sessions** to open the **Create Sessions** dialog box.
 - c. Click **Set Pairs** to open the **Set Pairs** dialog box.
 - d. Select a source volume, and a target volume, and click **Add** to make them a pair. Repeat this step as required.
 - e. Click **OK** to return to the **Create Sessions** dialog box.
5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To select a **Pairing Type**, expand the menu, and select one of the following. If you are not using the **Pairing Type** option, leave this field set to **None**.
 - ◆ **Use Exact Pairs** — Allows the system to pair up the volumes in the exact order that they were added to the group.
 - ◆ **Use Optimized Pairs** — Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.
2. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
3. Click **OK**.

Activating virtual copy sessions

Activating the copy session starts the copy-on-first-write mechanism and places the target volume in the Read/Write state. The target host can access the copy and has access to data on the source host until the copy session is terminated.

To activate virtual copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.

4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

Select a group, and click **Activate** to open the **Activate** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, and click **Activate** to open the **Activate** dialog box.

5. Select the type of **Source Volumes** to use.
6. Select the type of **Target Volumes** to use.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Establishing virtual copy sessions

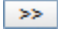
This procedure explains how to create and immediately activate clone copy sessions.

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more  , and select **Establish** to open the **Establish** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more  , and select **Establish** to open the **Establish** dialog box.

5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To select a **Pairing Type**, expand the menu, and select one of the following. If you are not using

the **Pairing Type** option, leave this field set to **None**.

- ◆ **Use Exact Pairs** — Allows the system to pair up the volumes in the exact order that they were added to the group.
 - ◆ **Use Optimized Pairs** — Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.
2. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
 3. Click **OK**.

Attaching and detaching preferred pairs

Pre-determining attached volume pairs eliminates the need to specify copy session target volumes from within the device group for create and activate operations.

The attached pairs will be used whenever a snap operation is requested for the specified device group. If a create or activate operation does not specify a volume pair from within the device group, the attached pair will automatically be used for the operation.


Before you begin:

You can only attach a standard source volume with a target VDEV. You cannot attach if the source volume is a BCV.

To attach/detach preferred pairs:

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Select a group, and click **View Details** to open its **Snap Sessions List** view.
5. Do the following depending on whether you are attaching or detaching volumes:

Attaching:

- a. Select one or more pairs, and click more  .
- a. Select **Attach** to open the **Attach** dialog box.
- b. To include BCVs, select the **BCV** option.
- c. Select the volumes to attach and click **OK**.

Detaching:

- a. Select one or more pairs and click **Detach**.
- a. Click **OK** in the confirmation message.

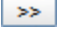
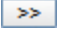
Duplicating virtual copy sessions

The duplicate TimeFinder/Snap feature allows you to duplicate a point-in-time copy of a virtual volume that is paired in a previously activated snap session to another virtual volume. This second point-in-time copy session actually resides with the source volume of the original snap session and is charged as part of the maximum number of sessions for that source volume. The duplicate snap is an actual copy of the virtual volume to another virtual volume.

Before you begin:

- ◆ Requires Engenuity 5875 or higher.
- ◆ Snap create and activate operations cannot be mixed between normal snap sessions and duplicate snap sessions within the same operation.
- ◆ The maximum number of duplicated sessions in the Created state is two.
- ◆ When a duplicate session is in the Created state, the original session cannot be terminated or recreated until the duplicate session is activated.

To duplicate virtual copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:
 - Group level:**
 - a. Select a group, click more , and select **Duplicate** to open the **Duplicate** dialog box.
 - Pair level:**
 - a. Select a group, and click **View Details** to open its **Details** view.
 - b. Select one or more pairs, click more , and select **Duplicate** to open the **Duplicate** dialog box.
5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Recreating virtual copy sessions

The snap recreate action allows you to recreate a snap session on an existing VDEV in order to prepare to activate a new point-in-time image.

Before you begin:

- ◆ This feature can only be used on sessions that have been previously activated.
- ◆ This feature requires a Symmetrix system running Enginuity 5874 or higher.
- ◆ For Symmetrix systems running Enginuity 5875 or higher, you can use this feature to recreate multivirtual snap sessions from thin and standard volumes.


To recreate virtual copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Recreate** to open the **Recreate** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more , select **Recreate** to open the **Recreate** dialog box.

5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of [options](#).
2. Click **OK**.

Restoring virtual copy sessions

The following types of restore operations can be performed for virtual copy sessions:

- ◆ Incremental restore back to the original source volume.

- ◆ Incremental restore to a BCV, which has been split from its original standard source volume but maintains the incremental relationship with the source.
- ◆ Full restore to any standard or split BCV outside of the existing copy session. The target volume of the restore must be of the same size and emulation type as the source volume.


Before you begin:

- ◆ With Engenuity 5875 or higher, you can use ORS control volumes as snap restore volumes when the volumes are in Push sessions and in the ORS Copied state.
- ◆ With Engenuity 5876 Q4 2012 and higher, you can perform a TimeFinder/Snap restore to a TimeFinder/Clone target. For example, volumes in an A > B > C cascaded session (where A > B is TimeFinder/Clone and B > C is TimeFinder/Snap) can copy data from volume C to volume A (via volume B). You can complete this operation without terminating the TimeFinder/Clone session, or any existing TimeFinder/Snap sessions off of the TimeFinder/Clone target. This feature is known as Persistent Restore to Target (PTT).


To restore virtual copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more  , and select **Restore** to open the **Restore** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more  , and select **Restore** to open the **Restore** dialog box.

5. Select the **Restore Type**.

Restore operations can be used to copy target data to another device (full restore), or back to the original source device (incremental restore). In the case of a full restore, the original session terminates and a copy session to the target of the restore starts. In the case of an incremental restore, the original session copy direction is reversed and changed data is copied from the target device to the source device. Restore operations require that the original session is differential and the source device is fully copied.

6. If performing a **Full** restore, click **Set Pairs** to open the **Set TimeFinder Snap Pairs** dialog box from which you can select the volumes to use in the operation.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu and select any number of

options.

2. Click **OK**.

Terminating virtual copy sessions

This procedure explains how to terminate an active virtual copy session at any time.

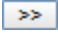
To terminate virtual copy sessions:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Terminate** to open the **Terminate** dialog box.

Pair level:

- a. Select a group, and click **View Details** to open its **Details** view.
- b. Select one or more pairs, click more , and select **Terminate Sessions** to open the **Terminate Sessions** dialog box.

5. *Optional:* Select **Use TGT Volumes** to use volumes from a local target list as targets.
6. *Optional:* If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of options.
2. Click **OK**.

Viewing Snap sessions

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Select a group, and click **View Details** to open its **Snap Sessions** list view.

Use the **Snap Sessions** list view to display and manage a group's snap pairs.

The following properties display:

- ◆ **Group Details** — Lists additional group details, including:
 - **Group Type** — The type of group. Possible values are: Regular, R1, R2, or R21.

- **Vender ID** — The company whose application created the group.
- **Group Valid** — Whether the group is valid (**Yes**) or invalid (**No**).
- **Group Creation Time** — Date/Time the group was created.
- **Application ID** — The application that created the group.
- **Group Modify Time** — Date/Time the group was last modified.
- ◆ **Session List** — Lists the group's snap sessions and their attributes, including:
 - **Source Volume** — The name of the source volume.
 - **Source LDev** — Logical name of the source volume.
 - **Target Volume** — The name of the target volume.
 - **Target LDev** — Theological name of the target volume.
 - **State** — The session state of the pair.
 - **Pool** — The name of the snap pool.
 - **Percent Copied** — The percentage of copying complete.

The following controls are available:

- ◆ **Create Sessions** — See [Creating virtual copy sessions](#) on page 333.
- ◆ **Activate** — See [Activating virtual copy sessions](#) on page 334.
- ◆ **Terminate** — See [Terminating virtual copy sessions](#) on previous page.
- ◆ **View Details** — [Viewing snap pair details](#)
- ◆ **Detach** — See [Attaching and detaching preferred pairs](#) on page 336.
- ◆ **Attach** — See [Attaching and detaching preferred pairs](#) on page 336.
- ◆ **Duplicate** — See [Duplicating virtual copy sessions](#) on page 337.
- ◆ **Establish** — See [Establishing virtual copy sessions](#) on page 335.
- ◆ **Recreate** — See [Recreating virtual copy sessions](#) on page 338.
- ◆ **Restore** — See [Restoring virtual copy sessions](#) on page 338.

Viewing snap pair details

1. Select the Symmetrix system.
2. Select **Data Protection** > **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Snap** to open the **TimeFinder Snap** view.
4. Select a group, and click **View Details** to open its **Snap Sessions** list view.
5. Select a pair and click **View Details** to open its **Details** view.
6. Use the snap session **Details** view to view and manage a snap session.

The following properties display:

- ◆ **Group** — Name of the snap source group.
- ◆ **Source Volume** — Name of the source volume.

- ◆ **Source LDev** — Logical name of the source volume.
- ◆ **Target Volume** — Name of the target volume.
- ◆ **Target LDev** — Logical name of the target volume.
- ◆ **Pair State** — Session state of the pair.
- ◆ **Number of Source Protected Tracks** — Number of protected tracks on the source.
- ◆ **SAVE Pool Name** — Name of the snap pool.
- ◆ **Percent Copied** — Percentage of copying complete.
- ◆ **Number of Source Modified Tracks** — Number of tracks modified on the source volume.
- ◆ **Number of Target Modified Tracks** — Number of tracks modified on the target volume.

The following controls are available:

- ◆ **Activate** — See [Activating virtual copy sessions](#) on page 334.
- ◆ **Terminate** — See [Terminating virtual copy sessions](#) on page 340.
- ◆ **Restore** — See [Restoring virtual copy sessions](#) on page 338.
- ◆ **Recreate** — See [Recreating virtual copy sessions](#) on page 338.
- ◆ **Detach** — See [Attaching and detaching preferred pairs](#) on page 336.
- ◆ **Attach** — See [Attaching and detaching preferred pairs](#) on page 336.
- ◆ **Establish** — See [Establishing virtual copy sessions](#) on page 335.

Snap session options

The following table describes the TimeFinder/Snap session options:

Table 8: TimeFinder/Clone session options

Session option	Description	Available with action
Consistent	Causes the source and VDEV pairs to be consistently activated.	Activate
Duplicate	Indicates that the action is being performed on a duplicate virtual copy session (that is, on a VDEV to a VDEV pair).	Create Activate Terminate
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.	Create Activate Terminate Restore Incremental Restore

Session option	Description	Available with action
Not Ready	Sets the VDEVs as Not Ready.	Activate Restore Incremental Restore
Restore	Must be used with the terminate action when terminating a restore session.	Terminate
Star	Indicates that the action is being performed on a volume that is in SRDF/Star mode.	Create Activate Recreate Terminate Restore
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.	Terminate

TimeFinder/Mirror

Establishing BCV pairs

The following explains how to establish standard volumes with BCV volumes:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Mirror** to open the **TimeFinder Mirror** view.
4. Do the following, depending on whether you want to perform this operation at the group level or the pair level:
 - Group level:**
 - a. Select a device group, and click **Establish** to open the **Establish - Device Group** dialog box.
 - Pair level:**
 - a. Select a device group, and click **View Details** to open the **Mirror Sessions List** view.
 - b. Select one or more pairs, and click **Establish** to open the **Establish Mirror Pair** dialog box.
5. Select a **Restore Type**:
 - ◆ **Incremental** — Copies to the BCV volume only the new data that was updated on the standard volume while the BCV pair was split.
 - ◆ **Full** — Copies the entire contents of the standard volume to the BCV volume.

6. If performing a full establish at the pair level, do the following:
 - a. Click **Set Pairs** to open the **Set TimeFinder Mirror Pairs** dialog box.
 - b. Select a **Source Volume** and a **Target Volume**, and click **Add** to make them a pair. Repeat this step as required.
 - c. Click **OK** to return to the **Establish - Mirror Pair** dialog box.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To specify **Full Establish Options**, expand the menu, and do the following:
 - a. Select one of the following options. If you are not using the **Full Establish Options**, leave this field set to **None**.
 - ◆ **Use Exact Pairs** — Allows the system to pair up the volumes in the exact order that they were added to the group.
 - ◆ **Use Optimized Pairs** — Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.
2. To attach **Session Options** to the operation, expand the menu, and select any number of options, as described in See [TimeFinder/Mirror session options](#) on page 347.
3. Click **OK**.

Restoring BCV pairs

The following explains how to copy data from the BCV volumes to the standard volumes:

1. Select the Symmetrix system.
2. Select **Data Protection** › **Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Mirror** to open the **TimeFinder Mirror** view.
4. Do the following, depending on whether you want to perform this operation at the group level or the pair level:

Group level:

 - a. Select a device group, and click **Restore** to open the **Restore - Device Group** dialog box.

Pair level:

 - a. Select a device group, and click **View Details** to open the **Mirror Sessions List** view.
 - b. Select one or more pairs, and click **Restore** to open the **Restore Mirror Pair** dialog box.
5. Select a **Restore Type**:
 - ◆ **Incremental** — Copies to the standard volume only the new data that was updated on the BCV volume while the BCV pair was split.
 - ◆ **Full** — Copies the entire contents of the BCV volume to the standard volume.
6. If performing a full establish at the pair level, do the following:
 - a. Click **Set Pairs** to open the **Set TimeFinder Mirror Pairs** dialog box.

- b. Select a **Source Volume** and a **Target Volume**, and click **Add** to make them a pair. Repeat this step as required.
 - c. Click **OK** to return to the **Restore - Mirror Pair** dialog box.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To specify **Full Restore Options**, expand the menu, and do the following:
 - a. Select one of the following options. If you are not using the **Full Establish Options**, leave this field set to **None**.
 - ◆ **Use Exact Pairs** — Allows the system to pair up the volumes in the exact order that they were added to the group.
 - ◆ **Use Optimized Pairs** — Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.
2. To attach **Session Options** to the operation, expand the menu, and select any number of options, as described in See [TimeFinder/Mirror session options](#) on page 347.
3. Click **OK**.

Splitting BCV pairs

The following explains how to split paired volumes to where each holds separate valid copies of the data:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Mirror** to open the **TimeFinder Mirror** view.
4. Do the following, depending on whether you want to perform this operation at the group level or the pair level:

Group level:

 - a. Select a device group, and click **Split** to open the **Split - Device Group** dialog box.

Pair level:

 - a. Select a device group, and click **View Details** to open the **Mirror Sessions List** view.
 - b. Select one or more pairs, and click **Split** to open the **Split Mirror Pair** dialog box.
5. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:


1. To attach **Session Options** to the operation, expand the menu, and select any number of options, as described in See [TimeFinder/Mirror session options](#) on page 347.
2. Click **OK**.

Cancelling BCV pairs


The following explains how to cancel a BCV pair relationship:

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Mirror** to open the **TimeFinder Mirror** view.
4. Do the following, depending on whether you want to perform this operation at the group level or the pair level:

Group level:

- a. Select a device group, click more , and select **Cancel** to open the **Cancel - Device Group** dialog box.

Pair level:

- a. Select a device group, and click **View Details** to open the **Mirror Sessions List** view.
 - b. Select one or more pairs, click more , and select **Cancel** to open the **Cancel Mirror Pair** dialog box.
5. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To attach **Session Options** to the operation, expand the menu, and select any number of options, as described in See [TimeFinder/Mirror session options](#) on the facing page.
2. Click **OK**.

Viewing mirror pairs

1. Select the Symmetrix system.
2. Select **Data Protection > Local Replication** to open the **Local Replication** dashboard.
3. From **Local Replication Type**, select **TimeFinder Mirror** to open the **TimeFinder Mirror** view.
4. Select a device group and click **View Details** to open the **Mirror Pairs** list view.
5. Use the **Mirror Pairs** list view to display and manage TimeFinder/Mirror sessions.

The following properties display:

- ◆ **Group Details** — Lists additional group details, including:
 - **Group Type** — The type of group. Possible values are: Regular, R1, R2, or R21.
 - **Vender ID** — The company whose application created the group.
 - **Group Creation Time** — Date/Time the group was created.
 - **Group Valid** — Whether the group is valid (**Yes**) or invalid (**No**).
 - **Group Modify Time** — Date/Time the group last modified.
 - **Application ID** — The application that created the group.

- ◆ **Session List** — Lists the group's clone pairs and their attributes, including:
 - **Source Volume** — The hexadecimal ID of the source volume.
 - **Source LDev** — The logical name of the source volume.
 - **Target Volume** — The hexadecimal ID of the target volume.
 - **Target LDev** — The logical name of the target volume.
 - **State** — The session state of the pair.
 - **CDGP** — Flags specific to the pair session in the form:
 - (C): X = The background copy setting is active for this pair.
 . = The background copy setting is not active for this pair.
 - (G): X = The Target volume is associated with a group.
 . = The Target volume is not associated with a group.
 - (D): X = The Clone session is a differential copy session.
 . = The Clone session is not a differential copy session.
 - (P): X = The precopy operation has completed one cycle.
 . = The precopy operation has not completed one cycle.
 - **Percent Copied** — The percentage of copying that is complete.

The following controls are available:

- ◆ **Establish** — See [TimeFinder/Mirror](#) on page 343.
- ◆ **Restore** — See [Restoring BCV pairs](#) on page 344.
- ◆ **Split** — See [Splitting BCV pairs](#) on page 345.
- ◆ **View Details** — See [Viewing mirror pair details](#) on page 1.
- ◆ **Cancel** — See [Cancelling BCV pairs](#) on previous page.

TimeFinder/Mirror session options

The following table describes the TimeFinder/Mirror session options:

Table 9: TimeFinder/Mirror session options

Session option	Description	Available with action
Bypass	Bypasses the Symmetrix system's exclusive locks for the local or remote array during mirror operations.	Split Full Restore Incremental Restore
Consistent	Causes the standard volumes being managed to be consistently split. Cannot be combined with the Instant option.	Split

Session option	Description	Available with action
Differential	Indicates that the split operation should initiate a differential data copy from the first mirror set member to the rest of the BCV mirror set members when the BCV pair split is done.	Split
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.	Full Establish Incremental Establish Split Full Restore Incremental Restore
Differential	Used with either the Copy or Precopy option to create an SDDF session for maintaining changed track information. This must be used when creating copy sessions on which you plan on issuing a Restore action.	Create Establish
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.	Create Establish Activate Restore Split Terminate
Not Ready	Sets the target volumes as Not Ready. Upon completion of a split action, the target volumes are set as Not Ready. When a restore is initiated, the standard volumes are set as Not Ready.	Split Full Restore Incremental Restore
Optimize	Optimizes volume pairings across the local Symmetrix array without regard for whether the volumes belong to different RDF (RA) groups. For remote volumes, use the Optimize Rag option.	Full Establish
Optimize Rag	Uses optimization rules to create remote BCV pairs from volumes	Full Establish

Session option	Description	Available with action
	within the same RDF (RA) group on a Symmetrix system.	
Protbcvest	Applies to two-way mirrored BCV volumes . Moves all mirrors of the BCV volume to join the mirrors of the standard volume.	Full Establish Incremental Establish
Protect	Indicates that the BCV should be write-protected before initiating a restore operation.	Split Full Restore Incremental Restore
Remote	Applicable only for split operations on a BCV RDF1 volume, or a restore operation from a BCV to a STD RDF2 volume. If this option is not specified, then the mode defaults to not propagate the data to the remote mirror of the RDF volume.	Split Full Restore Incremental Restore
Reverse	With a split operation, initiates a reverse data copy from one or more fixed BCV mirrors to the first (moving) mirror of the BCV upon the completion of the split operation. With an establish or restore operation, requests a verification check that the BCV's fixed mirror has valid data. If at establish or restore time you anticipate a need to perform future BCV reverse split operations, you must apply a reverse establish or restore so that no invalid tracks on the fixed BCV mirror become used.	Full Establish Incremental Establish Split Full Restore Incremental Restore
Star	Targets the action at volumes in SRDF/Star mode.	Full Establish Restore Split Cancel
SymForce	Forces an operation on the volume pair including pairs that would be rejected. Use caution when	Full Establish Incremental Establish Split

Session option	Description	Available with action
	checking this option because improper use may result in data loss.	Full Restore Incremental Restore

Device Groups

Creating device groups

1. Select the Symmetrix system.
2. Select **Data Protection > Device Group** to open the **Device Group** list view.
3. Click **Create** to open the **Create Device Group** wizard.
4. Type a **Device Group Name**.
5. Select a **Device Group Type**.
If the group type is **R1**, you can only add R1 volumes. If the group type is **R2**, you can add only R2 volumes. If the group type is **R21**, then you can only add R21 volumes.
6. Click **Next**.
7. Select the **Source** of the volumes to use when creating the group; either manual selection, or all the volumes in a storage group.
8. Do the following, depending on the source of the volumes:
 - Manual selection:**
 - a. Select the **Source Volume Type**.
 - b. Select one or more volumes and click **Add to Group**.
 - Storage group:**
Type or select the name of the **Storage Group**.
9. Click **Next**.
10. Select how to specify the **Target Volumes**, either manually or automatically.
11. Do the following, depending on how you are specifying the target volumes:
 - Automatically:**
 - a. Select whether to replicate the source volumes using **TimeFinder/Snap**, **TimeFinder Mirror**, or **TimeFinder/Clone**.
 - b. If you are replicating the source volumes with TimeFinder/Clone, select whether to add BCV or STD volumes to the device group. The volumes will be added with the TGT flag.
 - Manually:**
 - a. Click **Next**.
 - b. Select the **Target Volume Type**.
 - c. Select one or more volumes and click **Add to Group**.
12. Click **Next**.
13. Verify your selections in the **Summary** page. To change any of your selections, click **Back**. Note that some changes may require you to make additional changes to your configuration.
14. Click **Finish**.

Adding volumes to device groups


1. Select the Symmetrix system.
2. Select **Data Protection** > **Device Group** to open the **Device Group** list view.
You can also add volumes from the device group **Details** view, or from the list of volumes in the device group. From the volumes view click **Add** to add volumes.
3. Select the device group and click **Edit** to open the **Add Volumes to Device Group** view.
4. From the list of available volumes select the volumes and click **Add to Group**.
5. Click **OK**.

Removing volumes from device groups

1. Select the Symmetrix system.
2. Select **Data Protection** > **Device Group** to open the **Device Group** list view.
3. Select the device group and click **View Details** to open the **Details** view.
4. In the **Related Objects** panel, click a volume type to open the list view for the volume type.
5. Select one or more volumes and click **Remove**.
6. Click **OK** in the confirmation message.

Enabling device groups

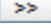
This procedure explains how to enable consistency protection for a device group consisting of SRDF/A volumes.

1. Select the Symmetrix system.
2. Select **Data Protection** > **Remote Replication** and expand the **Device Groups** folder.
3. Select a group, click more  , and select **Enable** to open the **Enable Device Group** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Disabling device groups

This procedure explains how to disable consistency protection for a device group consisting of SRDF/A volumes.

1. Select the Symmetrix system.
2. Select **Data Protection** > **Remote Replication** and expand the **Device Groups** folder.

3. Select a group, click ,more  , and select **Disable Consistency** to open the **Disable Device Group** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Renaming device groups

1. Select the Symmetrix system.
2. Select **Data Protection > Device Group** to open the **Device Group** list view.
3. Select the device group from the list and click **View Details** to open the **Details** view.
4. In the **Name** field, enter the new device group name and click **Apply** to save changes.

Deleting device groups

1. Select the Symmetrix system.
2. Select **Data Protection > Device Group** to open the **Device Group** list view.
3. Select the device group and click **Delete**.
4. Click **Delete** in the confirmation message.

Viewing device groups

1. Select the Symmetrix system.
2. Select **Data Protection > Device Group** to open the **Device Group** list view.

Use the **Device Group** list view to view and manage device groups.

The following properties display:

- ◆ **Name** — User-defined device group name.
- ◆ **Group Type** — Device configuration of the devices in the group. Possible values are: Regular, R1, R2, or R21.
- ◆ **Standards** — Number of standard devices in the device group.
- ◆ **BCVs** — Number of BCV devices in the device group.
- ◆ **VDEVs** — Number of virtual devices in the device group.
- ◆ **Targets** — Number of target devices in the device group.
- ◆ **Gatekeeper** — Number of gatekeeper devices in the device group.
- ◆ **Valid** — Indicates whether the device group is valid.

The following controls are available:

- ◆ **Create** — See [Creating device groups](#) on page 351.
- ◆ **View Details** — See [Viewing device group details](#) below.
- ◆ **Edit** — See [Adding volumes to device groups](#) on page 352.
- ◆ **Delete** — See [Deleting device groups](#) on previous page.

Viewing device group details

1. Select the Symmetrix system.
2. Select **Data Protection > Device Group** to open the **Device Group** list view.
3. Select the device group and click **View Details** to open its **Details** view.

The device group **Details** view allows you to view and manage a device group. This view contains [Properties](#), [Related Objects](#), and [Performance Views](#) panels.

Properties panel

The following properties display:

- ◆ **Name** — User-defined device group name.
- ◆ **Vendor ID** — System vendor ID.
- ◆ **Application ID** — Indicates which application created the device group.
- ◆ **Valid** — Indicates whether the device group is valid.

- ◆ **Device Group Priority QoS** — QoS priority setting for the device group.
- ◆ **Dynamic Cache Partition Name** — Dynamic cache partition to which the device group is assigned.
- ◆ **Dynamic Group Create Time** — Time the device group was created.
- ◆ **Dynamic Group Modify Time** — Time the device group was modified.
- ◆ **Symmetrix ID** — Symmetrix system serial number ID.
- ◆ **Remote Symmetrix ID** — Symmetrix system serial number ID of the remote Symmetrix system.
- ◆ **RBCV Symmetrix ID** — Symmetrix system serial number ID for the remote BCV.
- ◆ **RRBCV Symmetrix ID** — Symmetrix system serial number ID for remote RBCV.
- ◆ **Number of Associated Gatekeepers** — Number of gatekeeper devices in the device group.
- ◆ **Number of STD Devices in Group** — Number of standard devices in the device group.
- ◆ **Number of Locally-Associated BCVs** — Number of BCV devices associated with the device group.
- ◆ **Number of Locally-Associated VDEVs** — Number of virtual devices associated with the device group.
- ◆ **Number of Locally-Associated BCVs** — Number of local BCV devices associated with the device group.
- ◆ **Number of Remotely-Associated BCVs (STD RDF)** — Number of remote BCV devices associated with the device group.
- ◆ **Number of Remotely-Associated BCVs (BCV RDF)** — Number of BCV devices, associated with the device group, to be paired with remotely-attached BCV devices.
- ◆ **Number of Remotely-Associated RBCVs (RBCV RDF)** — Number of remote BCV devices associated with the device group.
- ◆ **Number of Remotely-Associated VDEVs** — Number of remote VDEV devices associated with the device group.
- ◆ **Number of Remotely-Associated TGTs** — Number of remote target devices associated with the device group.
- ◆ **Number of Hop2 BCVs (Remotely-associated Hop2 BCV)** — Number of BCVs on the second hop of the Cascaded SRDF configuration associated with the device group.
- ◆ **Number of Hop2 VDEVs (Remotely-associated Hop2 VDEV)** — Number of virtual devices on the second hop of the Cascaded SRDF configuration associated with the device group.
- ◆ **Number of Hop2 TGTs (Remotely-associated Hop2 TGT)** — Number of target devices on the second hop of the Cascaded SRDF configuration associated with the device group.

- ◆ **Number of Composite Groups** — Number of composite groups in the device group.
- ◆ **Pacing Capable** — Indicates if the device group allows write pacing capability.
- ◆ **Group Pacing**— Indicates if the device group is write pacing enabled or disabled.
- ◆ **Volume Pacing** — Indicates if the volumes in the device group are write pacing enabled or disabled.
- ◆ **Volume Write Pacing Exempt State** — Indicates if volume write pacing exemption capability is enabled or disabled.
- ◆ **Write Pacing Exempt Volumes** — Indicates if there are write pacing exempt volumes.
- ◆ **Effective Write Pacing Exempt State** — Indicates if effective write pacing exemption capability is enabled or disabled.

The following controls are available:

- ◆ **Create** — See [Creating device groups](#) on page 351.
- ◆ **Edit** — See [Adding volumes to device groups](#) on page 352.
- ◆ **Delete** — See [Deleting device groups](#) on page 353.
- ◆ **Apply** — Applies changes made in the **Properties** list. For example, renaming the storage group.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in or associated with the device group. Each group link is followed the name of the group, or by a number, indicating the number of objects in the corresponding view. For example, clicking **BCVs - 2** opens the view listing the two BCV devices contained in the device group.

Performance Views panel

The **Performance** panel links you to the performance analyze views for the group.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

Remote Replication

Managing remote replication sessions

The **Remote Replication** dashboard provides you with a single place to monitor and manage SRDF sessions on a Symmetrix system. This view contains the list of device groups that are valid for SRDF management. This includes device groups types R1, R2, and R21.

To access the **Remote Replication** dashboard:

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

The following properties display:

For the **Device Group** folder the total number of **Standard** and **BCV** volumes display for all listed devices groups.

For each device group the following properties display:

- ◆ **Source Group** — Device group name.
- ◆ **Standard** — Number of standard volumes.
- ◆ **BCV** — Number of BCV volumes.
- ◆ **State** — Current state of device group.
- ◆ **Group Type** — Device group type.
- ◆ **Group Valid** — Indicates if device group is valid or invalid for SRDF management.

The following controls are available:

- ◆ **Establish** — See [Establishing SRDF pairs](#) on page 368.
- ◆ **Split** — See [Splitting SRDF pairs](#) on page 377.
- ◆ **Suspend** — See [Suspending SRDF pairs](#) on page 378.
- ◆ **Restore** — See [Restoring SRDF pairs](#) on page 377.
- ◆ **View Details** — See [Viewing SRDF volume pairs](#) on page 362.
- ◆ **Set Mode** — See [Setting SRDF mode](#) on page 361.
- ◆ **Refresh** — See [Refreshing R1 or R2 volumes](#) on page 376.
- ◆ **RW Disable R2** — See [Read/write disabling R2 volumes](#) on page 373.
- ◆ **Write Disable** — See [Write disabling R1/R2 volumes](#) on page 375.
- ◆ **RW Enable** — See [Read/write enabling R1/R2 volumes](#) on page 373.
- ◆ **R1 Update** — See [Updating R1 volumes](#) on page 379.
- ◆ **Not Ready** — See [Making R1/R2 volumes not ready](#) on page 372.

- ◆ **Ready** — See [Making R1/R2 volumes ready](#) on page 371.
- ◆ **Deactivate SRDF/A** — See [Deactivating SRDF/A controls](#) on page 368.
- ◆ **Activate SRDF/A** — See [Controls](#) on page 367.
- ◆ **Invalidate** — See [Invalidating R1/R2 volumes](#) on page 371.
- ◆ **Move** — See [Moving SRDF pairs](#) on page 360.
- ◆ **Delete Pair** — See [Deleting SRDF pairs](#) on page 360.
- ◆ **Disable Consistency** — See [Disabling consistency protection](#) on page 400.
- ◆ **Enable Consistency** — [Enabling consistency protection](#) on page 400
- ◆ **Swap** — See [Swapping SRDF personalities](#) on page 379.
- ◆ **Failover** — [Failing over](#) on page 369
- ◆ **Failback** — [Failing back](#) on page 370
- ◆ **Resume** — [Resuming SRDF links](#) on page 374

SRDF

Configuration

Creating SRDF pairs

This procedure explains how to create SRDF pairs.

Before you begin:

Starting with Enginuity 5875, you can:

- ◆ Block the creation of an SRDF pair when the R2 is larger than the R1. This feature requires that you disable the SYMAPI_RDF_CREATEPAIR_LARGER_R2 option in the SYMAPI options file (enabled by default). For more information on disabling SYMAPI options, refer to the EMC Solutions Enabler Installation Guide.
- ◆ Create SRDF pairs containing standard and thin devices, or thin and diskless devices. To use this feature, the thin and diskless devices must be on a Symmetrix system running Enginuity 5875 or higher, and the standard device must be on a Symmetrix system running Enginuity 5875, 5773.50154, or 5671.

To create SRDF pairs:

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select the SRDF group and click **Create Pairs** to open the **Create SRDF Pairs** dialog box. This selection will determine the remote Symmetrix system.
4. Select the SRDF **Mirror Type** to add the volumes.
5. Select the **SRDF Mode** for the volume pairs. Depending on the existing configuration of the volumes in the group, not all modes may be available. For list of possible mode, refer to [SRDF group modes](#) on page 407.
6. *Optional:* Select the **Adaptive Copy Mode**. Possible values are:
 - ◆ **Enabled: Disk Mode** — When set, the Symmetrix system acknowledges all writes to source (R1) volumes as if they were local volumes. New data accumulates on the source (R1) volume and is marked by the source (R1) side as invalid tracks until it is subsequently transferred to the target (R2) volume. The remote director transfers each write to the target (R2) volume whenever link paths become available.
 - ◆ **Enabled: WP Mode** — When set, the Symmetrix 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.
7. Type the **Number of Volumes in the Range**.
8. Type or select the number of the first volume in the range of volumes on the local Symmetrix system (**Local Start Volume**).
9. Type or select the number of the first volume in the range of volumes on the remote Symmetrix system (**Remote Start Volume**).
10. *Optional:* Select to **Include Larger R2** volumes in the pair.
11. Click one of the following:
 - ◆ **OK** to create the pair now.
 - ◆ **Show Advanced** to set the advanced options, as described next.

Setting Advanced options:

1. Select any number of the following options:
 - ◆ **Invalidate** — Specifies the mirror to invalidate upon creation, if any. Possible values are:
 - **R1** — Invalidate the R1 mirror. This option is not supported on a volume pair with an R2 volume larger than the R1.
 - **R2** — Invalidates the R2 mirror. Starting with Enginuity 5875, this option is supported on a volume pair with an R2 volume larger than the R1.
 - **None** — Invalidates neither mirror.
 - ◆ **Establish** — Starts the RDF data copy upon creation.

- ◆ **Restore** — Resumes remote mirroring and initiates a data copy from the target (R2) side to the source (R1) side. Depending on the operation, all (full) or only changed track (incremental) are copied to the source (R1) side.
 - **Use RecoverPoint Tag** — Specifies that the operation will be performed on RecoverPoint volumes.
- ◆ **Format** — Specifies to clear all tracks on the R1 and R2 sides, ensuring no data exists on either side, and makes the R1 read write to the host
 - **Establish** — Starts the RDF data copy upon creation.
- ◆ **Consistency Exempt** — Allows you to move volumes to an RDF group supporting an active SRDF/A session without requiring other volumes in the group to be suspended.
- ◆ **Remote** — Requests a remote data copy flag.
- ◆ **No WD** — Specifies to not include volumes enabled for reads only.

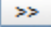
2. Click **OK**.

Deleting SRDF pairs

Deleting SRDF pairs cancels the dynamic SRDF pairing by removing the pairing information from the Symmetrix system and converting the volumes from SRDF to regular volumes. This operation can be performed on an SRDF group, a device group, or a composite group.

Half deleting SRDF pairs cancels the dynamic SRDF pairing information for one side (R1s or R2s) of the specified device pairs and converts the volumes from RDF to regular volumes. This operation can be performed on a device group or composite group.


To delete SRDF pairs:

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.
3. Select a group, click more , and select **Delete Pair** to open the **Delete Pair** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. Select the **Half Delete** option if deleting one side of the device pair.
6. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Moving SRDF pairs

The move action moves the SRDF pair from one SRDF group to another. The move type can be a full move or a half move. A half move specifies to move only the local half of the RDF pair. When using this action on an RDF 1 type pair, only the R1 volume is moved. When using this action on an RDF 2 type pair, only the R2 volume is moved.

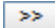
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

3. Select a group, click more , and select **Move** to open the **Move - Device Group** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. Select **New SRDF Group**.
6. Select **Move Type**.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

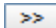
Setting SRDF mode

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.
Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Set Mode** to open the **Set Mode - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **SRDF Mode, Adaptive Copy Mode** and **AC Skew** to set the type of [SRDF session modes](#) (page 383).
- d. Select **Use Consistent** to set consistent transition from asynchronous to synchronous mode.
- e. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Set Mode** to open the **Set Mode - SRDF Pair** dialog box.
- c. Select **SRDF Mode, Adaptive Copy Mode** and **AC Skew** to set the type of [SRDF session modes](#) (page 383).
- d. Select **Use Consistent** to set consistent transition from asynchronous to synchronous mode.
- e. Click one of the following:
 - ◆ **OK** to start the operation now.

- ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Viewing SRDF volume pairs

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** to open the **Remote Replication** view.
3. Select a device group from the list and click **View Details** to open the **SRDF Session List** view. Use the **SRDF sessions List** view to display and manage SRDF group volume pairs and session status.
4. Click **Show Group Details** to display the device group details:
 - ◆ **Group Valid** — Indicates if device group is valid or invalid for SRDF management.
 - ◆ **Application ID** — Application name managing SRDF actions.
 - ◆ **Vendor ID** — Vendor name.
 - ◆ **Group Creation Time** — Group creation time stamp.
 - ◆ **Group Modify Time** — Group modification time stamp.
 - ◆ **Remote Symmetrix** — Remote Symmetrix ID
 - ◆ **Volume Pacing Exempt State** — Indicates if volume pacing exempt is enabled.
 - ◆ **Write Pacing Exempt State** — Indicates if write pacing exempt is enabled.
 - ◆ **Effective Write Pacing Exempt State** — Indicates if effective write pacing exempt is enabled.
5. Do one of the following to display **Local** or **Hop 2** SRDF link properties:

Local SRDF link properties:

Select **Local** for **SRDF Link Selection**.

The following properties display:

 - ◆ **Source Volume** — Source volume ID.
 - ◆ **Source LDev** — Source logical volume ID
 - ◆ **SRDF Group** — SRDF group ID.
 - ◆ **Remote Symmetrix** — Remote Symmetrix ID.
 - ◆ **Target Volume** — Target volume ID.
 - ◆ **State** — State of the RDF volume pairs.
 - ◆ **Volume State** — State of the source volume.
 - ◆ **Remote Volume State** — State of the remote volume.

Hop 2 SRDF link properties:

Select **Hop2** for **SRDF Link Selection**.

The following properties display:

- ◆ **Source LDev** — Source logical volume ID
- ◆ **Concurrent Volume** — Concurrent volume ID.
- ◆ **SRDF Group** — SRDF group ID.
- ◆ **Remote Symmetrix** — Remote Symmetrix ID.
- ◆ **Target Volume** — Target volume ID.
- ◆ **State** — State of the RDF volume pairs.
- ◆ **Volume State** — State of the source volume.
- ◆ **Remote Volume State** — State of the remote volume.

The following controls are available:

- ◆ **Establish** — See [Establishing SRDF pairs](#) on page 368.
- ◆ **Split** — See [Splitting SRDF pairs](#) on page 377.
- ◆ **Suspend** — [Suspending SRDF pairs](#) on page 378
- ◆ **Restore** — See [Restoring SRDF pairs](#) on page 377.
- ◆ **View Details** — See [Viewing SRDF volume pairs](#) on previous page.
- ◆ **Set Mode** — See [Setting SRDF mode](#) on page 361.
- ◆ **Refresh** — See [Refreshing R1 or R2 volumes](#) on page 376.
- ◆ **RW Disable R2** — See [Read/write disabling R2 volumes](#) on page 373.
- ◆ **Write Disable** — See [Write disabling R1/R2 volumes](#) on page 375.
- ◆ **RW Enable** — See [Read/write enabling R1/R2 volumes](#) on page 373.
- ◆ **R1 Update** — See [Updating R1 volumes](#) on page 379.
- ◆ **Not Ready** — See [Making R1/R2 volumes not ready](#) on page 372.
- ◆ **Ready** — See [Making R1/R2 volumes ready](#) on page 371.
- ◆ **Deactivate SRDF/A** — See [Deactivating SRDF/A controls](#) on page 368.
- ◆ **Activate SRDF/A** — See [Controls](#) on page 367.
- ◆ **Invalidate** — See [Invalidating R1/R2 volumes](#) on page 371.
- ◆ **Failover** — [Failing over](#) (page 369)
- ◆ **Failback** — [Failing back](#) (page 370)
- ◆ **Resume** — [Resuming SRDF links](#) (page 374)

Viewing SRDF volume pair details

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** to open the **Remote Replication** view.
3. Select a device group from the list and click **View Details** to open the **SRDF Session List** view.
4. Select a volume pair and click **View Details** to open the **SRDF Pair** details view.
5. Use the **SRDF Pair** details view and manage the SRDF group volume pair and session status.
This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **Device Group** — Device group ID.
- ◆ **Source Volume** — Source volume ID.
- ◆ **Source LDev** — Source logical device ID.
- ◆ **SRDF Group** — SRDF Group ID.
- ◆ **Remote Symmetrix** — Remote Symmetrix ID.
- ◆ **Remote SRDF Group** — Remote SRDF Group ID.
- ◆ **Target Volume** — Target volume ID.
- ◆ **Pair State** — Indicates volume pair state.
- ◆ **SRDF mode** — SRDF copy type.
- ◆ **Adaptive Copy Mode** — Indicates if adaptive copy mode is enabled.
- ◆ **Consistency State** — Indicates consistency state.
- ◆ **Consistency Exempt** — Indicates if consistency is exempt.
- ◆ **Link Status** — Indicates link state.
- ◆ **SRDF Domino** — Indicates SRDF Domino state.
- ◆ **SRDF Hop2 Group** — SRDF Hop2 Group ID.
- ◆ **Source Volume Invalid R1 Track Count** — Number of invalid R1 tracks on source volume.
- ◆ **Source Volume Invalid R2 Track Count** — Number of invalid R2 tracks on source volume.
- ◆ **Source Volume SRDF State** — Indicates source volume SRDF state.
- ◆ **Source Volume SRDF Type** — Indicates source volume SRDF type.
- ◆ **Source Volume Track Size** — Source volume track size.
- ◆ **Target Volume Invalid R1 Track Count** — Number of invalid R1 tracks on target volume.
- ◆ **Target Volume Invalid R2 Track Count** — Number of invalid R2 tracks on target volume.

- ◆ **Target Volume SRDF State** — Indicates target volume SRDF state.
- ◆ **Target Volume Track Size** — Target volume track size.
- ◆ **SRDF/A Pacing Capable** — Indicates if the SRDF pair allows write pacing capability.
- ◆ **Configured Group-level Exempt State** — Indicates if group-level write pacing exemption capability is enabled or disabled.
- ◆ **Effective Group-level Exempt State** — Indicates if effective group-level write pacing exemption capability is enabled or disabled.
- ◆ **Group Level Pacing State** — Indicates if group level write pacing is enabled or disabled.
- ◆ **Volume Level Pacing State** — Indicates if volume level write pacing is enabled or disabled.
- ◆ **SRDF/A Consistency Protection** — Indicates SRDF/A consistency protection state.
- ◆ **SRDF/A Average Cycle Time** — Average cycle time (seconds) configured for this session.
- ◆ **SRDF/A Minimum Cycle Time** — Minimum cycle time (seconds) configured for this session.
- ◆ **SRDF/A Cycle Number** — Indicates target volume SRDF state.
- ◆ **SRDF/A DSE Autostart** — Indicates DSE autostart state.
- ◆ **SRDF/A Session Number** — SRDF/A session number.
- ◆ **SRDF/A Session Priority** — Priority used to determine which SRDF/A sessions to drop if cache becomes full. Values range from 1 to 64, with 1 being the highest priority (last to be dropped).
- ◆ **SRDF/A Duration Of Last Cycle** — The cycle time (in secs) of the most recently completed cycle. It should be noted that in a regular case the cycles switch every ~30 sec, however, in most cases the collection interval is in minutes, which means some cycle times will be skipped. This an important counter to look at to figure out if SRDF/A is working as expected.
- ◆ **SRDF/A Flags** — RDFA Flags:

(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
- ◆ **SRDF/A Uncommitted Track Counts** — Number of uncommitted tracks.
- ◆ **SRDF/A Number of Volumes in Session** — Number of volumes in session.
- ◆ **SRDF/A Session Uncommitted Track Counts** — Number of uncommitted session tracks.
- ◆ **SRDF/A R1 DSE Used Track Count** — Number of tracks used for R1 DSE.
- ◆ **SRDF/A R1 Cache In Use Percent** — Percent of R1 cache used.

- ◆ **SRDF/A R1 Shared Track Count** — Number of R1 shared tracks.
- ◆ **SRDF/A R1 to R2 Lag Time** — Time that R2 is behind R1 (RPO). This is calculated as the last cycle time plus the time since last switch. In a regular case, the cycles switch every ~30 sec and the samples are taken every few minutes, therefore this counter may not show very significant data, however, when cycles elongate beyond the sample time, this counter can help indicate an estimate of the RPO.
- ◆ **SRDF/A R2 DSE Used Track Count** — Number of tracks used for R2 DSE.
- ◆ **SRDF/A R2 Cache In Use Percent** — Percent of R2 cache used.
- ◆ **SRDF/A Session Minimum Cycle Time** — Minimum cycle time (seconds) configured for this session.
- ◆ **SRDF/A Transmit Idle State** — Indicates SRDF/A transmit idle state.
- ◆ **SRDF/A Transmit Idle Time** — Time the transmit cycle has been idle.
- ◆ **Suspended State** — Suspended state.
- ◆ **Sqar Mode** — Indicates if SRDF pair is in a SQAR configuration.

The following controls are available:

- ◆ **Establish** — See [Establishing SRDF pairs](#) on page 368.
- ◆ **Split** — See [Splitting SRDF pairs](#) on page 377.
- ◆ **Suspend** — [Suspending SRDF pairs](#) on page 378
- ◆ **Restore** — See [Restoring SRDF pairs](#) on page 377.
- ◆ **View Details** — See [Viewing SRDF volume pairs](#) on page 362.
- ◆ **Set Mode** — See [Setting SRDF mode](#) on page 361.
- ◆ **Refresh** — See [Refreshing R1 or R2 volumes](#) on page 376.
- ◆ **RW Disable R2** — See [Read/write disabling R2 volumes](#) on page 373.
- ◆ **Write Disable** — See [Write disabling R1/R2 volumes](#) on page 375.
- ◆ **RW Enable** — See [Read/write enabling R1/R2 volumes](#) on page 373.
- ◆ **R1 Update** — See [Updating R1 volumes](#) on page 379.
- ◆ **Not Ready** — See [Making R1/R2 volumes not ready](#) on page 372.
- ◆ **Ready** — See [Making R1/R2 volumes ready](#) on page 371.
- ◆ **Deactivate SRDF/A** — See [Deactivating SRDF/A controls](#) on page 368.
- ◆ **Activate SRDF/A** — See [Controls](#) on the facing page.
- ◆ **Invalidate** — See [Invalidating R1/R2 volumes](#) on page 371.
- ◆ **Failover** — [Failing over](#) (page 369)

- ◆ **Failback** — [Failing back](#) (page 370)
- ◆ **Resume** — [Resuming SRDF links](#) (page 374)

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the SRDF group. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **SRDF Group** will open a view listing the two volumes contained in the SRDF group.

Controls


Activating SRDF/A controls

The activate action activates [SRDF/A control actions](#) (page 405) that detect cache overflow conditions and take corrective action to offload cache or slow down the host I/O rates to match the SRDF/A service rates.


1. Select the Symmetrix system.
2. Select **Data Protection** > **Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Activate SRDF/A** to open the **Activate SRDF/A - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Activate Type**.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Activate SRDF/A** to open the **Activate SRDF/A SRDF Pair** dialog box.
- c. Select **Activate Type**.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Deactivating SRDF/A controls

The deactivate action deactivates [SRDF/A control actions](#) (page 405) that detect cache overflow conditions and take corrective action.


1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Deactivate SRDF/A** to open the **Deactivate SRDF/A - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Deactivate Type**.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced options as described in [SRDF session options](#) on page 380. Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Deactivate SRDF/A** to open the **Deactivate SRDF/A SRDF Pair** dialog box.
- c. Select **Deactivate Type**.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced options as described in [SRDF session options](#) on page 380. Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Establishing SRDF pairs

This procedure explains how to start SRDF pair mirroring.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group and click **Establish** to open the **Establish - Device Group** dialog box.
- b. Select **Full** or **Incremental** session type.

- c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced options as described in [SRDF session options](#) on page 380. Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs and click **Establish** to open the **Establish Sessions** dialog box.
- c. Select **Full** or **Incremental** establish type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced options as described in [SRDF session options](#) on page 380. Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Failing over

In a period of scheduled downtime for maintenance, or after a serious system problem which has rendered either the host or Symmetrix system containing the source (R1) volumes unreachable, no read/write operations can occur on the source (R1) volumes. In this situation, the fail over operation should be initiated to make the target (R2) volumes read/write enabled to their local hosts.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.


Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more  and select **Failover** to open the **Failover** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select the fail over [SRDF session options](#) (page 380).
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.

- b. Select one or more pairs, click more  and select **Failover** to open the **Failback** dialog box.
- c. Select the fail over [SRDF session options](#) (page 380).
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Failing back

A fail back operation is performed when you are ready to resume normal SRDF operations by initiating read/write operations on the source (R1) volumes, and stopping read/write operations on the target (R2) volumes. The target (R2) volumes become read-only to their local hosts while the source (R1) volumes are read/write enabled to their local hosts.


1. Select the Symmetrix system.
2. Select **Data Protection** › **Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Failback** to open the **Failback** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select any of the available session options, as described in [SRDF session options](#) on page 380.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Failback** to open the **Failback** dialog box.
- c. Select any of the available session options, as described in [SRDF session options](#) on page 380.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Invalidating R1/R2 volumes

This procedure explains how to run internal checks to see if a volume swap is valid. To invoke this operation, the RDF pairs at the source must already be Suspended and Write Disabled or Not Ready.


1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Invalidate** to open the **Invalidate - Device Group** dialog box.
- b. Select **R1** or **R2** volume type.
- c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Invalidate** to open the **Invalidate SRDF Pair** dialog box.
- c. Select side **R1** or **R2**.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Making R1/R2 volumes ready

This procedure explains how to set the R1 (source) or R2 (target) volumes ready to their local hosts.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

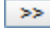
Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Ready** to open the **Ready - Device Group** dialog box.
- b. Select side **R1** or **R2**.

- c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Ready** to open the **Ready SRDF Pair** dialog box.
- c. Select **R1** or **R2** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

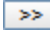
Making R1/R2 volumes not ready

This procedure explains how to set the source (R1) or the target (R2) volumes not ready to the local host.


1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Not Ready** to open the **Not Ready - Device Group** dialog box.
- b. Select side **R1** or **R2**.
- c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Not Ready** to open the **Not Ready SRDF Pair** dialog box.

- c. Select **R1** or **R2** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Read/write disabling R2 volumes

This procedure explains how to read/write disable target (R2) volumes to their local hosts.

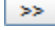
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **RW Disable R2** to open the **RW Disable R2 - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **RW Disable R2** to open the **RW Disable R2 SRDF Pair** dialog box.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Read/write enabling R1/R2 volumes

This procedure explains how to write enable the R1 (source) or R2 (target) volumes ready to their local hosts.

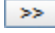
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more  , and select **RW Enable** to open the **RW Enable - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **RW Enable R1s** or **RW Enable R2s** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more  , and select **RW Enable** to open the **RW Enable - SRDF Pair** dialog box.
- c. Select **R1** or **R2** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.


Resuming SRDF links

This procedure explains how to resume I/O traffic on the SRDF links for all remotely mirrored RDF pairs in the group.

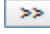
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more  , and select **Resume** to open the **Resume** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Resume** to open the **Resume** dialog box.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

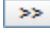
Write disabling R1/R2 volumes

This procedure explains how to write disable source (R1) volumes/target (R2) volumes to their local hosts.

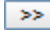
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Write Disable R2** to open the **Write Disable - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Write Disable R1s** or **Write Disable R2s** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Write Disable** to open the **Write Disable SRDF pair** dialog box.
- c. Select **R1** or **R2** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Refreshing R1 or R2 volumes

The refresh R1 action marks any changed tracks on the source (R1) volume to be refreshed from the target (R2) side. The Refresh R2 action marks any changed tracks on the target (R2) volume to be refreshed from the source (R1) side.

Before you begin:

To invoke this operation, the SRDF pair(s) must be in one of the following states:

- ◆ Suspended and Write Disabled at the source
- ◆ Suspended and Not Ready at the source
- ◆ Failed Over with the -force option specified


This operation is rejected if the target has invalid local (R2) tracks.

To refresh volumes:

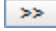
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more , and select **Refresh** to open the **Refresh - Device Group** dialog box.
- b. Select **R1** or **R2** volume type.
- c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **Refresh** to open the **Refresh SRDF Pair** dialog box.
- c. Select **R1** or **R2** volume type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Restoring SRDF pairs

This procedure explains how to restore data from the target (R2) volumes to the source (R1) volumes. When you fully restore SRDF pairs, the entire contents of the R2 volume is copied to the R1 volume. When you incrementally restore the R1 volume, only the new data that was changed on the R2 volume while the RDF group pair was split is copied to the R1 volume.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group and click **Restore** to open the **Restore - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Full** or **Incremental** restore type.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs and click **Restore** to open the **Restore Sessions** dialog box.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Splitting SRDF pairs

This procedure explains how to stop SRDF pair mirroring.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group and click **Split** to open the **Split - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Use Immediate** for immediate split on asynchronous devices.
- d. Click one of the following:

- ◆ **OK** to start the operation now.
- ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs and click **Split** to open the **Split** dialog box.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Suspending SRDF pairs

This procedure explains how to stop data transfer between SRDF pairs.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group and click **Suspend** to open the **Suspend - Device Group** dialog box.
- b. Select the **Hop2** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Use Immediate** or **Use Consistency Exempt** [SRDF session options](#) (page 380).
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Pair level:


- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs and click **Suspend** to open the **Suspend Sessions** dialog box.
- c. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

Swapping SRDF personalities

This procedure explains how to swap the SRDF volume designations for a specified device group. It changes source (R1) volumes to target (R2) volumes and target (R2) volumes to source (R1) volumes.

Half swapping SRDF personalities swaps one side of the RDF device designations for a specified group. It changes source (R1) volumes to target (R2) volumes or target (R2) volumes to a source (R1) volumes.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.
3. Select a group, click more  , and select **Swap** to open the **Swap** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. For optional **Refreshing R1 or R2 volumes** (page 376), select **R1, R2 or None**.
6. For option half swapping select **Half Swap**.
7. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced **SRDF session options** (page 380). Select the advanced options and click **OK**.

Updating R1 volumes


This procedure explains how to incrementally update R1 volumes with changed tracks from R2 volumes.

To update R1 volumes:

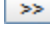
1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.

Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group, click more  , and select **R1 Update** to open the **R1 Update - Device Group** dialog box.
- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Select **Remote** if R1 volumes are a remote.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced **SRDF session options** (page 380). Select the advanced options and click **OK**.

Pair level:

- a. Select a group and click **View Details** to open the **SRDF pair list** view.
- b. Select one or more pairs, click more , and select **R1 Update** to open the **R1 Update SRDF Pair** dialog box.
- c. Select **Remote** if R1 volumes are a remote.
- d. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

This action can also be run from pair level details view. Select a pair and click **View Details**.

SRDF session options

Session option	Description	Available with action
Bypass	Bypasses the exclusive locks for the local and/or remote Symmetrix 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 Symmetrix systems.	Establish Failback Failover Restore Incremental Restore Split Suspend Swap Write Disable R1 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. To use this option on a composite group requires that the Symmetrix system be running Enginuity 5874 or higher.	Half Move Move Suspend

Session option	Description	Available with action
Establish	<p>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.</p> <p>To use this option on a composite group requires that the Symmetrix system be running Engenuity 5874 or higher.</p>	Failover
Force	<p>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.</p>	Establish Incremental Establish Restore Incremental Restore Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable Disable Swap
Immediate	<p>Causes the suspend, split, and failover actions on asynchronous volumes to happen immediately.</p>	Suspend Split Failover
NoWD	<p>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) volumes when used with the Invalidate R1 option and to the target (R2) volumes when used with the Invalidate R2 option.</p>	
SymForce	<p>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.</p>	Restore Incremental Restore Write Disable R1 Ready R1 Ready R2

Session option	Description	Available with action
		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
Star	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: <ul style="list-style-type: none"> ◆ Concurrent SRDF sites with SRDF/Star ◆ Cascaded SRDF sites with SRDF/Star <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>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

Session option	Description	Available with action
		Disable

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 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.
Synchronous	Provides the host access to the source (R1) volume on a write operation only after the Symmetrix system containing the target (R2) volume acknowledges that it has received and checked the data.
Asynchronous	<p>The Symmetrix 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 Symmetrix systems running Enginuity 5875, 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.

Migration

Understanding Virtual LUN Migration

Virtual LUN Migration (VLUN Migration) enables transparent, nondisruptive data mobility for both disk group provisioned and virtually provisioned Symmetrix 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 Symmetrix system. Migrations are performed while providing constant data availability and protection.

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 Symmetrix 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 Symmetrix system volume.

Virtual LUN Migration can be used to migrate regular Symmetrix 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.


Migrating regular storage group volumes

This procedure explains how to migrate all the regular volumes in a storage group.

Before you begin:

- ◆ Virtual LUN migration requires Engenuity 5874 or higher.
- ◆ To migrate volumes when access control is enabled, you must grant Symmetrix Optimizer access rights to the volumes.

To migrate regular storage group volumes:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** view.
3. Select a storage group, click more , and select **VLUN Migration** to open the **LUN Migration** dialog box.
4. Type a **Migration session name**. Migration session names must be less than 32 characters long and are case sensitive.
5. Select a **Target Disk Group**.
6. Select the RAID **Protection type**.
7. Select **Target type**. Choose **Create new volumes** to migrate to unconfigured volumes or **Use existing volumes** to migrate to configured volumes.
8. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
9. Click **OK** to create the migration session.

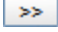
Migrating regular volumes

This procedure explains how to migrate individual regular volumes.

Before you begin:

- ◆ To migrate volumes when access control is enabled, you must grant Symmetrix Optimizer access rights to the volumes.
- ◆ Virtual LUN migration requires Engenuity 5874 or higher.

To migrate regular volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** view.
3. Expand the **Regular Volumes** folder, select the desired **Volume Type**, and click **View** to open the volume list.
4. Select one or more volumes, click more , and select **VLUN Migration** to open the **VLUN Migration** dialog box.
5. Type a **Migration session name**. Migration session names must be less than 32 characters and are case sensitive.
6. Select a **Target Disk Group**.
7. Select the RAID **Protection type**.

8. Select the **Target type**. Choose **Create new volumes** to migrate to unconfigured volumes or **Use existing volumes** to migrate to configured volumes.
9. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
10. Click **OK**.

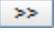
Migrating thin storage group volumes

This procedure explains how to migrate all the thin volumes in a storage group.

Before you begin:

- ◆ Virtual LUN migration requires Enginuity 5874 or higher.
- ◆ To migrate volumes when access controls is enabled, you must grant Symmetrix Optimizer access rights to the volumes.
- ◆ For instructions on migrating volumes on Symmetrix systems running Enginuity levels lower than 5874, see Migrating volumes.

To migrate thin storage group volumes:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** view.
3. Select a storage group, click more , and select **VLUN Migration** to open the **VLUN Migration** dialog box.
4. Type a **Migration session name**. Migration session names must be less than 32 characters and are case sensitive.
5. Select a **Target Pool**.
6. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
7. Click **OK**.


Migrating thin volumes

This procedure explains how to migrate individual thin volumes.

Before you begin:

- ◆ Virtual LUN migration requires Enginuity 5874 or higher.
- ◆ To migrate volumes when access controls is enabled, you must grant Symmetrix Optimizer access rights to the volumes.

To migrate selected thin volumes:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volumes** view.
3. Expand the **Virtual Volumes** folder, select **TDEV** Volume Type, and click **View** to open the **Thin Volumes** list.
4. Select one or more volumes, click more , and select **VLUN Migration** to open the **VLUN Migration** dialog box.
5. Type a **Migration session name**. The session name must be less than 32 characters long and is case sensitive.
6. Select a **Target Pool**.
7. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
8. Click **OK**.

Viewing VLUN migration sessions

1. Select the Symmetrix system.
2. Select **Data Protection > Migration** to open the **Migration** sessions list view.

Use the this view to display and manage migration sessions.

The following properties display:

- ◆ **Name** — Migration session name.
- ◆ **Invalid Tracks** — Number of invalid tracks for the volume pair.
- ◆ **Status** — Migration session status.
- ◆ **Percentage** — Percentage of the session completed.

The following controls are available:

- ◆ **Terminate** — [Terminating a VLUN migration session](#) on page 251
- ◆ **View Details** — [Viewing VLUN migration session details](#) below

Viewing VLUN migration session details

1. Select the Symmetrix system.
2. Select **Data Protection > Migration** to open the **Migration** sessions list view.
3. Select a session and click **View Details** to open its **Details** view.

Use this view to display details on a migration session. This view contains two panels, [Properties](#) and [Source and Target Information](#).

Properties panel

The following properties display:

- ◆ **Name** — Migration session name.
- ◆ **Invalid Tracks** — Number of invalid tracks for the volume pair.
- ◆ **Status** — Migration session status.
- ◆ **Percentage** — Percentage of the session completed.
- ◆ **Target Type** — Type of target volume.
- ◆ **Thin Pool** — If the target type is thin, this is the name of the pool containing the thin volume.

Source and Target Information

The following properties display:

- ◆ **Source** — Source volumes in the migration session.
- ◆ **Target** — Target volumes in the migration session.
- ◆ **Number of Target Volumes** — Number of target volumes in the session.
- ◆ **Invalid Tracks** — Number of invalid tracks for the volume pairs in the session.
- ◆ **Status** — Migration session status for the pair.

Replication Groups and Pools

SRDF/A DSE Pools

Creating SRDF/A DSE pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Click **Create** to open the **Create DSE Pool** dialog box. You can also create DSE pools from the pool details view.
4. Type a **Pool Name**. DSE pool names can contain up to 12 alpha-numeric characters. The only special character allowed is the underscore (_). The name DEFAULT_POOL is reserved for SAVE volumes that are enabled and not in any other pool.
5. Select the volumes to add and click one of the following:
 - ◆ **OK** to perform the operation now.
 - ◆ **Show Advanced** to set the advanced options, as described next.

Setting Advanced options:

1. To enable the new pool members when creating the pool, select **Enable New Pool Member**.
2. Click **OK**.

Deleting SRDF/A DSE pools


1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select a pool and click **Delete**.
4. Click **Delete** in the confirmation message.

Adding volumes to SRDF/A DSE pools


1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select a pool and click **Add** to open the **Add Volumes to DSE Pool** dialog box.
4. Select the volumes to add and click **Add to Pool**.
5. Click **OK**.

Enabling all volumes in SRDF/A DSE pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.

3. Select a pool and click more , and select **Enable All**.
4. Click **OK** in the confirmation message.

Disabling all volumes in SRDF/A DSE pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select a pool, click more , and select **Disable All**.
4. Click **OK** in the confirmation message.

Viewing SRDF/A DSE pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.

Use this list view to display and manage the SRDF/A DSE pools on a Symmetrix system.

The following properties display:

- ◆ **Name** — Name of the pool.
- ◆ **Configuration** — Configuration of the volumes in the pool.
- ◆ **Technology** — Technology on which the volumes in the pool reside.
- ◆ **Emulation** — Emulation type.
- ◆ **Pool State** — Whether the pool is Enabled or Disabled.
- ◆ **% Used** — Percent of pool used.
- ◆ **Used (MB)** — Total used space in MB.
- ◆ **Free (MB)** — Total free space in MB.

The following controls are available:

- ◆ **Create** — See [Replication Groups and Pools](#) on previous page.
- ◆ **Add** — See [Adding volumes to SRDF/A DSE pools](#) on previous page.
- ◆ **View Details** — See [Viewing SRDF DSE pool details](#) on the facing page.
- ◆ **Delete** — See [Deleting SRDF/A DSE pools](#) on previous page.
- ◆ **Enable All** — See [Enabling all volumes in SRDF/A DSE pools](#) on previous page.
- ◆ **Disable All** — See [Disabling all volumes in SRDF/A DSE pools](#) above.

Viewing SRDF DSE pool details

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select the pool and click **View Details** to open its **Details** view.

Use the SRDF/A DSE Pool **Details** view to display and manage a TimeFinder/Snap pool. It contains [Properties Performance Views](#), and [Related Objects](#) panels.

Properties panel

The following properties display:

- ◆ **Symmetrix ID** — Symmetrix system on which the pool resides.
- ◆ **Pool Name** — Name of the pool.
- ◆ **Type** — Pool type.
- ◆ **Emulation** — Emulation type.
- ◆ **RAID Protection** — Protection level of the volumes in the pool.
- ◆ **Number of Volumes** — Number of volumes in the pool.
- ◆ **Disabled Volumes** — Number of disabled volumes in the pool.
- ◆ **Enabled Volumes** — Number of enabled volumes in the pool.
- ◆ **Capacity** — Sum of all enabled and disabled volumes in the pool.
- ◆ **Enabled Capacity** — Sum of all enabled volumes in the pool.
- ◆ **Free Capacity** — Total free space in MB.
- ◆ **Technology** — Technology on which the volumes in the pool reside.
- ◆ **State** — Whether the pool is Enabled or Disabled.

The following controls are available:

- ◆ **Create** — See [Replication Groups and Pools](#) on page 389.
- ◆ **Add** — See [Adding volumes to SRDF/A DSE pools](#) on page 389.
- ◆ **Apply** — Applies changes made in the **Properties** list.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Object panel

The **Related Objects** panel provides links to views for objects contained in and associated with the pool. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **SAVE Volumes - 2** will open a view listing the two SAVE volumes contained in the pool.

Performance Views panel

The **Performance** panel links you to the performance analyze views for the policy.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

TimeFinder Snap Pools

Creating TimeFinder/Snap pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Click Create to open the **Create Snap Pool** dialog box.
4. Type a **Pool Name**. Snap pool names can contain up to 12 alpha-numeric characters. The only special character allowed is the underscore (_). The name DEFAULT_POOL is reserved for SAVE volumes that are enabled and not in any other pool.
5. Select one or more volumes and click either of the following:
 - ◆ **OK** — To add the selected volumes.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

To enable new volumes in the pool, select **Enable New Pool Member**.

The Total Enabled Pool Capacity in GB is displayed.

Adding volumes to TimeFinder/Snap pools

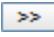
1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a pool and click **Add** to open the **Add Volumes to Snap Pool** dialog box.
4. Select one or more volumes and click either of the following:
 - ◆ **OK** — To add the selected volumes.
 - ◆ **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:


To enable new volumes in the pool, select **Enable New Pool Member**.

The Total Enabled Pool Capacity in GB displays.

Enabling all volumes in TimeFinder/Snap pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a snap pool, click more  , and select **Enable All**.
4. Click **OK** in the confirmation message.

Disabling all volumes in TimeFinder/Snap pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a snap pool, click more  , and select **Disable All**.
4. Click **OK** in the confirmation message.

Deleting TimeFinder/Snap Pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a pool and click **Delete**.
4. Click **Delete** in the confirmation message.

Viewing TimeFinder/Snap pools

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.

Use the **TimeFinder Snap Pools** list view to display and manage the TimeFinder/Snap pools on a Symmetrix system.

The following properties display:

- ◆ **Name** — Name of the pool.
- ◆ **Configuration** — Configuration of the volumes in the pool.
- ◆ **Technology** — Technology on which the volumes in the pool reside.
- ◆ **Emulation** — Emulation type.
- ◆ **Pool State** — Whether the pool is Enabled or Disabled.
- ◆ **% Used** — Percent of pool used.
- ◆ **Used (MB)** — Total used space in MB.
- ◆ **Free (MB)** — Total free space in MB.

The following controls are available:

- ◆ **Create** — See [TimeFinder Snap Pools](#) on page 392.
- ◆ **Add** — See [Adding volumes to TimeFinder/Snap pools](#) on page 392..
- ◆ **View Details** — See [Viewing TimeFinder/Snap pool details](#) below.
- ◆ **Delete** — See [Deleting TimeFinder/Snap Pools](#) on previous page.
- ◆ **Enable All** — See [Enabling all volumes in TimeFinder/Snap pools](#) on previous page.
- ◆ **Disable All** — See [Disabling all volumes in TimeFinder/Snap pools](#) on previous page.

Viewing TimeFinder/Snap pool details

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups and Pools > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a pool and click **View Details** to open its **Details** view.

The **Snap Pool Details** view allows you to view to manage a TimeFinder/Snap pool. It contains, [Properties](#), [Performance Views](#), and [Related Objects](#) panels.

Properties panel

The following properties display:

- ◆ **Symmetrix ID** — Symmetrix system on which the pool resides.
- ◆ **Pool Name** — Name of the pool.
- ◆ **Type** — Pool type.
- ◆ **Emulation** — Emulation type.
- ◆ **RAID Protection** — Protection level of the volumes in the pool.
- ◆ **Number of Volumes** — Number of volumes in the pool.
- ◆ **Disabled Volumes** — Number of disabled volumes in the pool.
- ◆ **Enabled Volumes** — Number of enabled volumes in the pool.
- ◆ **Capacity** — Sum of all enabled and disabled volumes in the pool.
- ◆ **Enabled Capacity** — Sum of all enabled volumes in the pool.
- ◆ **Free Capacity** — Total free space in MB.
- ◆ **Technology** — Technology on which the volumes in the pool reside.
- ◆ **State** — Whether the pool is Enabled or Disabled.

The following controls are available:

- ◆ **Create** — See [TimeFinder Snap Pools](#) on page 392..
- ◆ **Add** — See [Adding volumes to TimeFinder/Snap pools](#) on page 392.
- ◆ **Delete** — See [Deleting TimeFinder/Snap Pools](#) on page 393.
- ◆ **Enable** — See [Enabling all volumes in TimeFinder/Snap pools](#) on page 393.
- ◆ **Disable** — See [Disabling all volumes in TimeFinder/Snap pools](#) on page 393.
- ◆ **Apply** — Applies changes made in the **Properties** list.
- ◆ **Cancel** — Cancels changes made in the **Properties** list.

Related Object panel

The **Related Objects** panel provides links to views for objects contained in and associated with the pool. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **SAVE Volumes - 2** will open a view listing the two SAVE volumes contained in the storage group.

Performance Views panel

The **Performance** panel links you to the performance monitor and analyze views for the snap pool.

This panel only displays when the Performance option is installed. This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

SRDF Groups

Creating SRDF groups

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

Before you begin:

- ◆ The maximum number of supported RDF groups differs by Enginuity version:

Enginuity	Maximum number of RDF Groups supported		Group numbers
	per Symmetrix system	per director	
5874 or higher	250	64	1 to 250

- ◆ If the RDF interaction is between Enginuity 5874 (or higher) and Enginuity pre-5874, the maximum RDF group number of RDF is determined by the lower Enginuity level, however, the number of RDF groups per Symmetrix system and number of RDF groups defined on a director is determined solely by the Enginuity level running on the Symmetrix system where the group is being defined.

To create an SRDF group:

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Click **Create** to open the **Create SRDF group** dialog box.
4. Select a **Communication Protocol** to use when moving data across the SRDF links. The value you select here will populate the **Director** list.
5. Type a **SRDF group Label** (name).
6. Type an **SRDF group Number**.
7. Select the local **Director** through which the group will communicate.
8. Select the **Remote Symmetrix ID**.
9. Type a **Remote SRDF group Number**.
10. Select the remote **Director** through which the group will communicate.
11. Click one of the following:
 - ◆ **OK**
 - ◆ **Show Advanced** to set the advanced options, as described next.

Setting Advanced options:

1. Select a **Link Limbo Period**. This is a length of time for the Symmetrix system to continue checking the local SRDF link status. (The range is 0-120 seconds, default is 10.) If the link status is Not Ready after the link limbo time, the volumes are made Not Ready to the link.
2. Select (enable) **Link Domino** for the local group. With this feature enabled from either the local or remote side of group's RDF links, failure of the group's last remaining link will make all source (R1) volumes in the group unavailable (not ready) to their host when an R1-side operation occurs. This ensures that the data on the source (R1) and target (R2) devices is always in synch.
3. Select (enable) **Auto Link Recovery** for the local group. With this feature enabled, once the link failure is corrected, volumes that were ready to their host before the failure will automatically be restored to the ready state.
4. Select (enable) **Software Compression** for the local group. This enables SRDF software data compression for SRDF groups defined on GigE, or Fibre Channel. Although you can enable/disable software compression on the R2 side, the setting of hardware compression on the R1 side is what enables or disables the feature. This feature requires Enginuity 5874 or higher.
5. Select (enable) **Hardware Compression** for the local group. This enables SRDF hardware data compression on an SRDF group defined on a GigE director. Although you can enable/disable hardware compression on the R2 side, the setting of hardware compression on the R1 side is what enables or disables the feature. This feature requires Enginuity 5875 or higher.
6. Select a **Remote Link Limbo Period**. This is a length of time for the Symmetrix system to continue checking the remote SRDF link status. (The range is 0-120 seconds, default is 10.) If the link status is Not Ready after the link limbo time, the volumes are made Not Ready to the link.


7. Select (enable) **Remote Link Domino** for the remote group. With this feature enabled from either the local or remote side of group's RDF links, failure of the group's last remaining link will make all source (R1) volumes in the group unavailable (not ready) to their host when an R1-side operation occurs. This ensures that the data on the source (R1) and target (R2) volumes is always in synch.
8. Select (enable) **Remote Auto Link Recovery** for the remote group. With this feature enabled, once the link failure is corrected, volumes that were ready to their host before the failure will automatically be restored to the ready state.
9. Click **OK**.

Modifying SRDF groups


1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select a group and click **Edit** to open the **Modify SRDF group** dialog box.
4. Do any number of the following steps:
 - a. Select the local **Director** — through which the group will communicate.
 - b. Select the **Remote Director** through which the group will communicate.
 - c. Select a **Link Limbo Period**. The length of time for the Symmetrix system to continue checking the local SRDF link status. (The range is 0-120 seconds, default is 10.) If the link status is Not Ready after the link limbo time, the volumes are made Not Ready to the link.
 - d. Select (enable) **Link Domino** for the local group. With this feature enabled from either the local or remote side of group's RDF links, failure of the group's last remaining link will make all source (R1) volumes in the group unavailable (not ready) to their host when an R1-side operation occurs. This ensures that the data on the source (R1) and target (R2) devices is always in synch.
 - e. Select (enable) **Auto Link Recovery** for the local group. With this feature enabled, once the link failure is corrected, volumes that were ready to their host before the failure will automatically be restored to the ready state.
 - f. Select (enable) **Software Compression** for the local group. This enables SRDF software data compression for SRDF groups defined on GigE, or Fibre Channel. Although you can enable/disable software compression on the R2 side, the setting of hardware compression on the R1 side is what enables or disables the feature. This feature requires Enginuity 5874 or higher.
 - g. Select (enable) **Hardware Compression** for the local group. This enables SRDF hardware data compression on an SRDF group defined on a GigE director. Although you can enable/disable hardware compression on the R2 side, the setting of hardware compression on the R1 side is what enables or disables the feature. This feature requires Enginuity 5875 or higher.
 - h. Select a **Remote Link Limbo Period**. This is a length of time for the Symmetrix system to continue checking the remote SRDF link status. (The range is 0-120 seconds, default is 10.) If the link status is Not Ready after the link limbo time, the volumes are made Not Ready to the link.

- i. Select (enable) **Remote Link Domino** for the remote group. With this feature enabled from either the local or remote side of group's RDF links, failure of the group's last remaining link will make all source (R1) volumes in the group unavailable (not ready) to their host when an R1-side operation occurs. This ensures that the data on the source (R1) and target (R2) volumes is always in synch.
 - j. Select (enable) **Remote Auto Link Recovery** for the remote group. With this feature enabled, once the link failure is corrected, volumes that were ready to their host before the failure will automatically be restored to the ready state.
5. Click **OK**.

Setting SRDF/A group attributes

1. Select a Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select a group, click more , and select **SRDF/A Settings** to open the **Set SRDF/A Settings** dialog box.
4. Type the **RDFA Min(imum) Cycle Time**. This is the minimum amount of time (in seconds) the Symmetrix will wait before attempting to perform an RDF/A cycle switch. Possible values range from 5 to 59 seconds.
5. Type the **RDFA Session Priority**. This priority is used to determine which RDF/A session to drop if cache is full. Possible values range from 1 (highest) to 64 (lowest).
6. *Optional:* Enable **Transmit Idle Time** to preserve the data in cache (if the link is idle) and then retry transmitting the data. This option must be enabled on both local and remote sides.
7. Click **OK**.

Setting SRDF/A DSE attributes


1. Select a Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select a group, click more , and select **SRDF/A DSE** to open the **Set SRDF/A DSE Settings** dialog box.
4. Select the pool.
5. Type the percentage of the Symmetrix system's write pending limit (**Threshold**). Once the cache usage of all active groups in the Symmetrix system exceeds this limit, data tracks for this group start to spill over to disks. Possible values are from 20 to 100, with 50 being the default.
6. *Optional:* Select (enable) the SRDF/A write pacing feature to automatically start for the group when an SRDF/A session is activated (**Autostart**). This feature must be activated for host write I/O pacing to be invoked.

7. Manually **Activate/Deactivate** the SRDF/A Delta Set Extension (DSE) feature. DSE allows SRDF/A cache to be extended by offloading some or all of the session cycle data to preconfigured disks or pools. Possible values are:
 - ◆ **No change** — Leaves the current write pacing setting.
 - ◆ **Activate** — Activates the feature for the local side of the SRDF link.
 - ◆ **Activate Both Sides** — Activates the feature for both sides of the SRDF link.
 - ◆ **Deactivate** — Deactivates the feature for the local side of the SRDF link.
 - ◆ **Deactivate Both Sides** — Deactivates the feature for both sides of the SRDF link.

Starting with Engenuity 5773.150, this feature is supported with thin devices.

8. Click **OK**.

Setting SRDF/A pace attributes

1. Select a Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select a group, click more  , and select **SRDF/A Pace** to open the **Set SRDF/A Pace Settings** dialog box.
4. Type the maximum I/O delay to apply to each host write I/O when the pacing algorithm is invoked (**Pacing Delay**). Possible values range from 1 to 1,000,000 usec (0.000001 to 1 second), with 50,000 (0.05 seconds or 50 ms) being the default.
5. Type the minimum cache percentage when host write pacing will start (**Threshold**). Possible values range from 1 to 99, with 60% being the default.
6. *Optional:* Select to set the threshold on both the R1 and R2 sides (**Both Sides**).
7. *Optional:* Set the following write pacing attributes for the RDF group, the volumes in the group, or both:
 - a. Select (enable) the SRDF/A write pacing feature to automatically start when an SRDF/A session is activated (**Autostart**). This feature must be activated for host write I/O pacing to be invoked.
 - b. Manually **Activate/Deactivate** the SRDF/A write pacing feature for the RDF group. Setting this option to **No Change** leaves the current write pacing setting.

SRDF/A write pacing can only be activated when the SRDF/A session is active.
8. Click **OK**.


Swapping SRDF groups

When you swap the SRDF personality of the designated SRDF devices, the source (R1) devices become target (R2) devices and the target (R2) devices become source (R1) devices.

Before you begin:


- ◆ Swapping or half swapping SRDF pairs in composite groups requires that the Symmetrix system be running Enginuity 5874 or higher.
- ◆ Starting with Enginuity 5875, swaps on SRDF pairs in a split state are allowed (composite groups only).

To swap SRDF groups:

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select an SRDF group, click more , and select **Swap** to open the **Swap SRDF group** dialog box.
4. Select the mirror to refresh (**Refresh R1, Refresh R2**).
5. Optional: Select to Start refreshing the selected mirror (**Start Copy**).
6. If the RDF group you are swapping is part of an SRDF/Star configuration, select **Star**.
7. Click one of the following:
 - ◆ **Add to Job List** to add this task to the job list, from which you can schedule or run the task at your convenience. For more information, refer to [Scheduling jobs](#) on page 100 and [Previewing/Running jobs](#) on page 99.
 - ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Enabling consistency protection

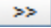
This procedure explains how to enable consistency protection for a device group consisting of SRDF/A volumes.

1. Select the Symmetrix system.
2. Select **Data Protection > Remote Replication** and expand the **Device Groups** folder.
3. Select a group, click more  and select **Enable** to open the **Enable Device Group** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [session options](#). Select the advanced options and click **OK**.


Disabling consistency protection

This procedure explains how to disable consistency protection for a device group consisting of SRDF/A volumes.

1. Select the Symmetrix system.

2. Select **Data Protection** › **Remote Replication** and expand the **Device Groups** folder.
3. Select a group, click more  , and select **Disable** to open the **Disable Device Group** dialog box.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. Click one of the following:
 - ◆ **OK** to start the operation now.
 - ◆ **Show Advanced** to set the advanced [SRDF session options](#) (page 380). Select the advanced options and click **OK**.

Deleting SRDF groups

1. Select the Symmetrix system.
2. Select **Data Protection** › **Replication Groups & Pools** › **SRDF groups** to open the **SRDF groups** list view.
3. Select an SRDF group, click more  , and select **Delete**.
4. Click **OK** in the confirmation message.

Viewing SRDF groups

1. Select the Symmetrix system.
2. Select **Data Protection** › **Replication Groups & Pools** › **SRDF groups** to open the **SRDF groups** list view.

Use the **SRDF groups** list view to display and manage SRDF groups.

The following properties display

- ◆ **SRDF Group** — RDF group number.
- ◆ **SRDF Group Label** — RDF group label.
- ◆ **Remote SRDF group** — Remote RDF group number.
- ◆ **Remote** — Remote Symmetrix serial ID.
- ◆ **SRDF group Flags** — SRDF group flags.
- ◆ **Volume Count** — Number of volumes in the group.
- ◆ **Copy Jobs** — Maximum number of RDF copy jobs per RDF group.
- ◆ **Link Limbo (sec)** — Number of seconds (0-10) for the Symmetrix system to continue checking the local RDF link status.
- ◆ **CSRMTDA** — See [RDF flags](#) on page 406.
- ◆ **Minimum Cycle** — Minimum time to wait before attempting an SRDF/A cycle switch. Values range from 5 to 59 seconds.
- ◆ **Session Priority** — Priority used to determine which SRDF/A sessions to drop if cache becomes full. Values range from 1 to 64, with 1 being the highest priority (last to be dropped).
- ◆ **Transmit Idle** — Time the transmit cycle has been idle.

The following controls are available:

- ◆ **Create** — See [SRDF Groups](#) on page 395.
- ◆ **Create Pairs** — See [SRDF](#) on page 358.
- ◆ **Edit** — See [Modifying SRDF groups](#) on page 397.
- ◆ **View Details** — See [Viewing SRDF group details](#) on the facing page.
- ◆ **Assign Dynamic Cache Partition** — See [Assigning volumes to dynamic cache partitions](#) on page 131.
- ◆ **Delete** — See [Deleting SRDF groups](#) on previous page.
- ◆ **Swap Groups** — See [Swapping SRDF groups](#) on page 399.
- ◆ **SRDF/A DSE Settings** — See [Setting SRDF/A DSE attributes](#) on page 398.
- ◆ **SRDF/A Pacing Settings** — See [Setting SRDF/A pace attributes](#) on page 399.
- ◆ **SRDF/A Settings** — See [Setting SRDF/A group attributes](#) on page 398.

Viewing SRDF group details

1. Select the Symmetrix system.
2. Select **Data Protection > Replication Groups & Pools > SRDF groups** to open the **SRDF groups** list view.
3. Select the SRDF group and click **View Details** to open its **Details** view.

Use the **SRDF Group Detail** view to view the properties of an SRDF group. This view contains two panels, [Properties](#) and [Related Objects](#).

Properties panel

The following properties display:

- ◆ **RDF (RA) Group Number** — RDF group number.
- ◆ **RDF (RA) Group Label** — RDF Group Label.
- ◆ **Remote RDF (RA) Group Number(s)** — Remote RA group number(s).
- ◆ **Remote Symmetrix ID(s)** — Remote Symmetrix serial ID(s).
- ◆ **Remote RDF (RA) Director Number(s)** — Remote RA director numbers.
- ◆ **Remote Director Ident(s)** — Remote director identifier(s).
- ◆ **RDF (RA) Group Flags** — SRDF group flags.
- ◆ **Prevent Auto Link Recovery** — Indicates the state of preventing automatic data copy across RDF links upon recovery.
- ◆ **Copy Jobs** — Maximum number of RDF copy jobs per RDF group.
- ◆ **Prevent RAs Online Upon Power On** — Indicates the state of preventing the RDF. directors from automatically coming back online with power on.
- ◆ **Link Domino** — Sets the domino mode for the source (R1) volumes.
- ◆ **Link Config** — Link configuration.
- ◆ **Director Config** — Indicates the Fibre adapter type.
- ◆ **RDF (RA) Group Configuration** — RA group configuration.
- ◆ **Farpoint** — RA group is configured for FarPoint.
- ◆ **RDFA Flags** — RDFA Flags:

(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

- ◆ **Link Limbo (sec)** — Number of seconds (0-10) for the Symmetrix system to continue checking the local RDF link status.
- ◆ **Minimum Cycle Time** — Minimum cycle time (seconds) configured for this session.
- ◆ **Session Priority** — Priority used to determine which SRDF/A sessions to drop if cache becomes full. Values range from 1 to 64, with 1 being the highest priority (last to be dropped).
- ◆ **Transmit Idle Time** — Time the transmit cycle has been idle.
- ◆ **Transmit Idle** — Whether SRDF/A Transmit Idle state is active for the RDF group.
- ◆ **R1 Side Percent Cache in Use** — Percent of system write pending cache slots used by the R1 side.
- ◆ **R2 Side Percent Cache in Use** — Percent of system write pending cache slots used by the R2 side.
- ◆ **Dynamic Cache Partition Name** — Cache partition name.
- ◆ **SRDF/A DSE Status** — Indicates if SRDF/A DSE is active.
- ◆ **SRDF/A DSE Autostart** — Indicates if SRDF/A DSE is automatically enabled when an SRDF/A session is activated for the group.
- ◆ **SRDF/A DSE Threshold** — Percentage of the Symmetrix systems write pending limit.
- ◆ **SRDF/A Write Pacing Status** — Indicates if SRDF/A write pacing is active.
- ◆ **SRDF/A Contains Not Pace Capable Volume** — Indicates if group contains write pacing capable volumes.
- ◆ **SRDF/A Write Pacing Delay** — Max delay allowed for host I/O in seconds.
- ◆ **SRDF/A Write Pacing Threshold** — Minimum cache percentage when host write pacing will start.
- ◆ **SRDF/A Write Pacing Autostart** — Indicates if the SRDF/A write pacing feature is automatically activated when an SRDF/A session is activated.
- ◆ **SRDF/A Write Pacing Supported** — Indicates if SRDF/A write pacing is supported
- ◆ **RDF Software Compression** — Indicates if software compression is enabled/disabled on the RDF group.
- ◆ **RDF Single Round Trip** — Indicates if single round trip is enabled/disabled on the RDF group.
- ◆ **RDF Hardware Compression** — Indicates if hardware compression is enabled/disabled on the RDF group.
- ◆ **Device Pacing Supported** — Indicates if SRDF/A Device Write Pacing is supported.
- ◆ **Device Pacing Activated** — Group-level pacing status of the SRDF/A session. The status of the feature can be Active, Inactive, N/A.
- ◆ **Device Pacing Autostart** — Indicates if SRDF/A Device Write Pacing autostart is enabled.

- ◆ **SRDF Software Compression** — Indicates if SRDF software compression is enabled or disabled.
- ◆ **SRDF Single Round Trip** — Indicates if SRDF single round trip is enabled.
- ◆ **SRDF Hardware Compression** — Indicates if SRDF hardware compression is enabled or disabled.
- ◆ **SRDF Software Compression Support** — Indicates if SRDF software compression is supported on the Symmetrix system.
- ◆ **SRDF Group Software Compression** — Indicates if SRDF software compression is enabled or disabled for the SRDF group.
- ◆ **SRDF Hardware Compression Support** — Indicates if SRDF hardware compression is supported on the Symmetrix system.
- ◆ **SRDF Group Hardware Compression** — Indicates if SRDF hardware compression is enabled or disabled for the SRDF group.
- ◆ **SQAR Mode** — Indicates if SRDF group is in a SQAR configuration.

The following controls are available:

- ◆ **Create Pairs** — See [SRDF](#) on page 358.
- ◆ **Edit** — See [Modifying SRDF groups](#) on page 397.
- ◆ **Assign Dynamic Cache Partition** — See [Assigning volumes to dynamic cache partitions](#) on page 131.
- ◆ **Delete** — See [Deleting SRDF groups](#) on page 401.
- ◆ **Swap Groups** — See [Swapping SRDF groups](#) on page 399.
- ◆ **SRDF/A DSE Settings** — See [Setting SRDF/A DSE attributes](#) on page 398.
- ◆ **SRDF/A Pacing Settings** — See [Setting SRDF/A pace attributes](#) on page 399.
- ◆ **SRDF/A Settings** — See [Setting SRDF/A group attributes](#) on page 398.

Related Objects panel

The **Related Objects** panel provides links to views for objects contained in and associated with the SRDF group. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Volumes - 2** will open a view listing the two volumes contained in the SRDF group.

SRDF/A control actions

These actions are invoked to detect cache overflow conditions and take corrective action to offload cache or slow down the host I/O rates to match the SRDF/A service rates.

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.
	Write Pacing This feature extends the availability of SRDF/A by preventing conditions that result in cache overflow on both the R1 and R2 sides.	Group write pacing Group level write pacing is supported on Symmetrix systems running Enginuity 5874.207.166 and higher.	Activates SRDF/A write pacing at the group level.
		Group & Volume Write Pacing	Activates SRDF/A write pacing at the group level and the volume level
		Volume Write Pacing Volume write pacing is supported on Symmetrix systems running Enginuity 5875 and higher.	Activates SRDF/A write pacing at the volume level.
Write Pacing Exempt	N/A	Activates write pacing exempt. Write pacing exempt allows you to remove a volume from write pacing.	

RDFA flags

Flag	Status
(C)onsistency	X = Enabled . = Disabled - = N/A
(S)tatus	A = Active I = Inactive - = N/A

Flag	Status
(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

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 Symmetrix system containing the target (R2) volume acknowledges that it has received and checked the data.
- ◆ **Asynchronous** — The Symmetrix 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 Symmetrix systems running Enginuity 5875, you can put an RDF relationship into Asynchronous mode when the R2 volume is a snap source volume.

- ◆ **Semi Synchronous** — The Symmetrix 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 Symmetrix system containing the target (R2) volume checks and acknowledges receipt of each write.
- ◆ **AC WP Mode On** — (adaptive copy write pending) the Symmetrix 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 Symmetrix 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) Off** — Turns off the (R2 NR_If_Invalid) On mode.

RecoverPoint

Understanding RecoverPoint

EMC 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.

For Symmetrix systems running Enginuity 5874 Q2 2001 SR or higher, Unisphere for VMAX allows you to tag/untag volumes for use with RecoverPoint.

Tagging and untagging volumes for RecoverPoint

This procedure explains how to tag (enable) and untag (disable) volumes for RecoverPoint. Enabling volumes makes them accessible to the RecoverPoint Appliance.

You can tag/untag volumes for RecoverPoint at the volume or storage group level.

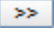
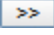
Before you begin:

Volumes that are part of an RDF pair cannot be tagged for RecoverPoint.

RecoverPoint operations on Unisphere for VMAX require Engenuity 5875 Q2 2011 SR or higher on the Symmetrix system.

To tag/untag volumes at the volume level:



From the Volumes dashboard:

1. Select the Symmetrix system.
2. Select **Storage > Volumes** to open the **Volume Dashboard**.
3. In the **Volume Type** panel, double-click the type of volume you want to tag/untag. The **Volumes** list view opens.
4. To tag volumes, select volumes, click more , and select **Tag for RecoverPoint**.
To untag volumes, select volumes, click more , and select **Untag for RecoverPoint**.
5. Click **OK** in the confirmation message.

Untagging from the RecoverPoint tagged volumes view:

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Recover Point Volumes** to open the **RecoverPoint Volumes** view.
3. Single or multi-select (hold shift key and select) volumes and click **Untag for RecoverPoint**.
4. Click **OK** in the confirmation message.

To tag/untag volumes at the storage group level:

1. Select the Symmetrix system.
2. Select **Storage > Storage Groups** to open the **Storage Groups** list view.
3. To tag the storage group, select it, click more , and select **Tag for RecoverPoint**.
To untag the storage group, select it, click more , and select **Untag for RecoverPoint**.
4. Click **OK** in the confirmation message.

Viewing RecoverPoint sessions

1. Select a Symmetrix system.
2. Select **Data Protection > OpenReplicator > RecoverPoint Sessions** to open the **RecoverPoint Sessions** list view.

Use the **RecoverPoint Sessions** list view to view RecoverPoint sessions on the Symmetrix system.

The following properties display:

- ◆ **Cluster name** — Session name.
- ◆ **Control volume**—Control volume name.
- ◆ **Remote volume** — Remote volume name.
- ◆ **Status** — Session status.
- ◆ **Protected Tracks** — Number of protected tracks.

The following controls are available:

- ◆ **View Details** — See [Viewing RecoverPoint session details](#) below.

Viewing RecoverPoint session details

1. Select a Symmetrix system.
2. Select **Data Protection > OpenReplicator > RecoverPoint Sessions** to open the **RecoverPoint Sessions** list view.

Select a session and click **View Details** to open the session details view.

Use the **RecoverPoint Sessions** details view to view session details.

The following properties display:

- ◆ **Cluster Name** — Session name.
- ◆ **Control Volume** — Control volume name.
- ◆ **Remote Volume**— Remote volume name.
- ◆ **Remote Volume Specification** — Indicates the remote volume name format.
- ◆ **Status** — Session status.
- ◆ **Copy pace** — Copy pace value.
- ◆ **Number of Protected Tracks** — Number of protected tracks.

Viewing RecoverPoint tagged volumes

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Recover Point Volumes** to open the **RecoverPoint Volumes** view.

The following properties display:

- ◆ **Name** — Volume name.
- ◆ **Type** — Volume volume.
- ◆ **Status** — Volume status.
- ◆ **Reserved** — Indicates if volume is reserved.
- ◆ **Capacity (GB)** — Volume capacity in GB.
- ◆ **Emulation** — Volume emulation type.

The following control is available:

- ◆ **View Details** — View details on the selected sessions.
- ◆ **Untag for RecoverPoint** — See [Viewing RecoverPoint tagged volume details](#) below.

Viewing RecoverPoint tagged volume details

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Recover Point Volumes** to open the **RecoverPoint Volumes** view.
3. Select the volume and click **View Details** to open its **Details** view.

This view allows you to view the volume details.

The following properties display:

- ◆ **Name** — Volume name.
- ◆ **Physical Name** — Physical name.
- ◆ **Volume Identifier** — Volume identifier.
- ◆ **Type** — Volume configuration.
- ◆ **Encapsulated Volume** — Whether the volumes is encapsulated. Relevant for external disks only.
- ◆ **Encapsulated WWN** — World Wide Name for encapsulated volume. Relevant for external disks only.
- ◆ **Status** — Volume status.
- ◆ **Reserved** — Whether the volume is reserved.
- ◆ **Capacity (GB)** —Volume capacity in GBs.

- ◆ **Capacity (MB)** — Volume capacity in MBs.
- ◆ **Capacity (Cylinders)** —Volume capacity in cylinders.
- ◆ **Emulation** — Volume emulation.
- ◆ **Symmetrix ID** — Symmetrix system on which the volume resides.
- ◆ **Symmetrix Volume ID** — Symmetrix volume name/number.
- ◆ **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- ◆ **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- ◆ **Nice Name** — Nice name generated by Symmetrix Enginuity.
- ◆ **WWN** — World Wide Name of the volume.
- ◆ **DG Name** — Name of the device group in which the volume resides, if applicable.
- ◆ **CG Name** — Name of the device group in which the volume resides, if applicable.
- ◆ **Attached BCV** — Defines the attached BCV to be paired with the standard volume.
- ◆ **Attached VDEV TGT Volume** — Volume to which this source volume would be paired.
- ◆ **RDF Type** — RDF configuration.
- ◆ **Geometry - Type** — Method used to define the volume's geometry.
- ◆ **Geometry - Number of cylinders** — Number of cylinders, as defined by the volume's geometry.
- ◆ **Geometry - Sectors per Track** — Number of sectors per track, as defined by the volume's geometry.
- ◆ **Geometry - Tracks per Cylinder** — Number of tracks per cylinder, as defined by the volume's geometry.
- ◆ **Geometry - 512 Block Bytes** — Number of 512 blocks, as defined by the volume's geometry.
- ◆ **Geometry Capacity (GB)** — Geometry capacity in GBs.
- ◆ **Geometry Limited** — Indicates whether an encapsulated volume has a Symmetrix cylinder size larger than the reported user-defined geometry.
- ◆ **SSID** — Subsystem ID.
- ◆ **Capacity (Tracks)** — Capacity in tracks.
- ◆ **SA Status** — Volume SA status.
- ◆ **Host Access Mode** — Host access mode.
- ◆ **Pinned** —Whether the volume is pinned.

- ◆ **RecoverPoint Tagged** — Whether or not the volume is tagged for RecoverPoint.
- ◆ **Service State** — Service state.
- ◆ **Defined Label Type** — Type of user-defined label.
- ◆ **Dynamic RDF Capability** — RDF capability of the volume.
- ◆ **Mirror Set Type** — Mirror set for the volume and the volume characteristic of the mirror.
- ◆ **Mirror Set DA Status** — Volume status information for each member in the mirror set.
- ◆ **Mirror Set Invalid Tracks** — Number of invalid tracks for each mirror in the mirror set.
- ◆ **Priority QoS** — Priority value assigned to the volume. Valid values are 1 (highest) through 16 (the lowest).
- ◆ **Dynamic Cache Partition Name** — Name of the cache partition.
- ◆ **XtremSW Cache Attached** — Whether the volume is currently controlled by cache cards.

The following controls are available:

- ◆ **Untag for RecoverPoint** - See [Tagging and untagging volumes for RecoverPoint](#) on page 410.

Open Replicator

Creating Open Replicator copy sessions

Before you begin:

There are many rules and limitations for running Open Replicator sessions. Refer to the *Solutions Enabler Symmetrix Symmetrix Migration CLI Product Guide* before creating a session.

For quick reference refer to [Open Replicator session options](#) on page 421.

To create a copy session:

1. Select the Symmetrix system.
2. Select **Data Protection** > **Open Replicator** > **Open Replicator SAN View** to open the **Open Replicator SAN View**.
3. Click **Create Copy Session** to open the **Create Copy Session** wizard.
4. Select a **Copy Direction** and **Copy Operation**.
5. Click **Next**.

The **Source - Remote Volumes** lists the remote volumes from the Open Replicator remote volumes list view. The **Target - Control Volumes** lists all the control volumes that can be paired with the remote volumes.

For a cold push session, one control volume can concurrently push data to up to 16 remote volumes. For cold pull, hot push, and hot pull sessions only one control volume can push/pull to one remote device.

6. Select a remote volume and target volume, then click **Add Pair**.
If the pair is valid it is added to the **Volume Pairs** list.
7. Click **Remove Pair** to edit the **Volume Pairs** list.
8. Click **Next**.
9. Enter **Session Name**.
10. Enter **Copy Pace** value (0 - slowest to 9 - fastest).

With offline copying, there is a slight pause between each track write. You can speed up a copy operation by reducing or eliminating this pause. While in the CopyInProgress or CopyOnAccess state, set a pace value higher than the default of 5. Setting the copy pace to 9 eliminates this pause..

11. Select the [Open Replicator session options](#) (page 421) and click **Next**.
12. View session **Summary** and click **Finish** to create session or click **Back** to edit session options.

Managing Open Replicator sessions

1. Select the Symmetrix system.

2. Select **Data Protection > Open Replicator > Open Replicator SAN View** to open the **Open Replicator SAN View**.

Use this view to view select remote volumes to use for Open Replicator and FLM (Federated Live Migration) copy sessions. The list of volumes can be filtered further by selecting the objects within the tree views.

Open Replicator view tree view lists

The following properties display:

- ◆ **Control Ports** — Control system director (Name and port number).
- ◆ **Remote Ports** — Remote system ports (Port WWN, Array ID, and Vendor).
- ◆ **Remote Volumes** — Remote system volumes (Volume WWN and capacity (GB)).

Filtering remote volumes for sessions

To filter the remote volumes, single or multi-select (hold shift key and select) the items in the tree view. As each selection is made, the filtered results table of remote volumes is updated to reflect the current filter criteria. The filtered remote volumes list is used in the create session wizard to determine which control volumes that can be paired with the remote volumes.

Filtered results table

The following properties display:

- ◆ **Remote volume WWN** — Remote system volume WWN.
- ◆ **Vendor** — System vendor.
- ◆ **Capacity (GB)** — Capacity, in GB, of the volume.
- ◆ **Volume** — Volume name.
- ◆ **LUN** — LUN ID.
- ◆ **Emulation** — Volume emulation.
- ◆ **Thin** — Indicates if the volume is a thin volume.
- ◆ **In Session** — Indicates if the session is in process.
- ◆ **State** — Volume state
- ◆ **Reserved** — Indicates if the volume is reserved.

The following controls are available:

- ◆ **Create Copy Session** — See [Creating Open Replicator copy sessions](#) on previous page.
- ◆ **Create FLM Session** — See [Creating a FLM session](#) on page 428.
- ◆ **Rescan** — Refreshes the remote volumes list.

Activating Open Replicator sessions

Before you begin:

The copy session must be in a created or recreated state before you can activate it.

To activate a session:

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Activate** to open **Activate Session** dialog box.
4. Select a copy option. Refer to [Open Replicator session options](#) on page 421 for session copy and control options.
5. Click **OK**.

Restoring Open Replicator sessions

Before you begin:

The restore operation restores the copy session back to the control volume by pulling back only the changed tracks from the remote volume. The session must have been created with differential copying, and must be in the copied state. Hot or cold differential push sessions can be restored.

To restore a session:

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Restore** to open **Restore Session** dialog box.
4. Select any number of the available options. Refer to [Open Replicator session options](#) on page 421 for session control options.
5. Click **OK**.

Renaming Open Replicator sessions

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Rename** to open **Rename Session** dialog box.
4. Type a new name for the session.
5. Click **OK**.

Removing Open Replicator sessions

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Remove** to open **Remove Session** dialog box, and click **OK**.
An error message will display if the session is in a state that does not allow the session to be removed.

Setting Open Replicator session background copy mode

This procedure sets the session background copy mode for an ORS session that has already been created.

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Set Mode** to open **Set Mode** dialog box.
4. Select the background copy mode. Refer to [Open Replicator session options](#) on page 421 for session control options.
5. Click **OK**.

Setting Open Replicator session donor update off

This procedure deactivates donor update for a session that was created with donor update.

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Donor Update Off** to open **Set Donor Update Off** dialog box.
4. Select the [Open Replicator session options](#) on page 421.
5. Click **OK**.

Setting Open Replicator session front end zero detection off

This procedure deactivates front end zero detection for a session that was created with front end zero.

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.

3. Select a session and click **Frontend Zero Off** to open **Set Frontend Zero Off** dialog box. Refer to [Open Replicator session options](#) on page 421 for session control options.
4. Click **OK**.

Setting Open Replicator session pace

This procedure sets how fast data copies between volumes during an ORS session. Values can range from 0 to 9, with 0 being the fastest pace, and 9 being the slowest pace. If set to 0, there is no inserted delay time and the replication will proceed as fast as possible.

Values of 1 - 9 add delays, which takes longer to complete copying but conserves system resources. The default for both online (hot) replication and offline (cold) replication is 5.

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions** to open the **Open Replicator Sessions** list view.
3. Select a session and click **Set Pace** to open **Set Pace** dialog box.
4. Type a **Pace** value (0 - fastest to 9 - slowest).
5. Click **OK**.

Setting Open Replicator ceiling

The Open Replicator ceiling value is the percentage of bandwidth available for background copying. You should only set this value after understanding the bandwidth being used by other applications. By default, the ceiling value is NONE.

1. Select a Symmetrix system.
2. Select **System > Dashboard > Front End Directors** to open the Front End Directors list view.
3. Select a director and click **Set ORS Ceiling** to open the **Set ORS Ceiling** dialog box.
4. Type a **Open Replicator Ceiling** value from 1 (minimum) to 100 (maximum) and click **OK**.

Terminating Open Replicator sessions

1. Select the Symmetrix system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions View** to open the **Open Replicator SAN View**.
3. Select a session and click **Terminate** to open the **Terminate** confirmation dialog box.
4. Select terminate options. Refer to [Open Replicator session options](#) on page 421 for session control options.
5. Click **OK**.


Viewing Open Replicator sessions

1. Select the Symmetrix system.

2. Select **Data Protection > Open Replicator > Open Replicator Sessions View** to open the **Open Replicator Sessions View**.

Use the this view to view and manage Open Replicator sessions.

The following properties display:

- ◆ **Session** — ORS session name.
- ◆ **Control Volume** — Control volume name.
- ◆ **Remote Volume** — Remote volume name.
- ◆ **Status** — Session status.
- ◆ **Protected Tracks** — Number of protected tracks.
- ◆ See [Open Replicator session options](#) on the facing page. indicated as enabled for the session by  :
 - **FLM session**
 - **Background Copy**
 - **Differential Copy**
 - **Pull Session**
 - **Cold Copy Session**
 - **Donor Update**
 - **Front End Zero Detection**

The following controls are available:

- ◆ **View Details** — See [Viewing Open Replicator session details](#) below.
- ◆ **Activate** — See [Activating Open Replicator sessions](#) on page 417.
- ◆ **Terminate** — See [Terminating Open Replicator sessions](#) on previous page.
- ◆ **Remove** — See [Removing Open Replicator sessions](#) on page 418.
- ◆ **Front End Zero off** — See [Setting Open Replicator session donor update off](#) on page 418.
- ◆ **Donor Update Off** — See [Setting Open Replicator session donor update off](#) on page 418.
- ◆ **Rename** — See [Renaming Open Replicator sessions](#) on page 417.
- ◆ **Remove** — See [Removing Open Replicator sessions](#) on page 418.
- ◆ **Restore** — See [Restoring Open Replicator sessions](#) on page 417.
- ◆ **Set Pace** — See [Setting Open Replicator session pace](#) on previous page.
- ◆ **Set Mode** — See [Setting Open Replicator session background copy mode](#) on page 418.

Viewing Open Replicator session details

1. Select the Symmetrix system.

2. Select **Data Protection > Open Replicator > Open Replicator Sessions View** to open the **Open Replicator Sessions View**.

Select a session and click **View Details** to open the session details view.

The following properties display:

- ◆ **Session** — ORS session name.
- ◆ **Control Volume** — Control volume name.
- ◆ **Remote Volume** — Remote volume name.
- ◆ **Remote Volume Specification** — Remote volume specification.
- ◆ **Status** — Session status.
- ◆ **Percent Complete** — Percent tracks copied.
- ◆ **Copy Pace** — Copy Pace value (0 - slowest to 9 - fastest, default is 5).
- ◆ **Number of Protected Tracks** — Number of protected tracks.
- ◆ **Number of Modified Tracks** — Number of modified tracks.
- ◆ **Background Copy** — Indicates if background copying is enabled.
- ◆ **Differential Copy** — Indicates if differential copying is enabled.
- ◆ **Pull Session** — Indicates if session is a pull session = Yes, or a push session = No.
- ◆ **Cold Copy Session** — Indicates if session is a cold copy session = Yes, or a hot copy session = No.
- ◆ **Donor Update** — Indicates if donor update is enabled.
- ◆ **FLM Session** — Indicates if session is a FLM session.
- ◆ **RecoverPoint Session** — Indicates if session is a RecoverPoint session.
- ◆ **Standard ORS Session** — Indicates if session is a standard session.
- ◆ **Front-End Zero Detection** — Indicates if front-end zero detection is enabled.

Open Replicator session options

Session Option	Used with Command	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.
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 Force 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 Force 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 Force_Copy option.
Frontend_zero	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.

Session Option	Used with Command	Description
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 CopyInProgress.
Pull		A pull operation copies data to the control device from the remote device.
Push		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.

Open Replicator flags

Flag	Status
C Background copying	X = Enabled . = Disabled
D Differential copying	X = Enabled . = Disabled
S Copy direction	X = Pushing data to the remote volume(s) . = Pulling data from the remote volume(s)
H Copy operation	X = Hot copy session . = Cold copy session
U Donor update	X = Enabled . = Disabled
T Session type	M = Migration session (FLM). R = RecoverPoint session. S = Standard ORS session.

Flag	Status
Z Front-end zero detection	X =Enabled . = Disabled
*	Failed session can be reactivated.

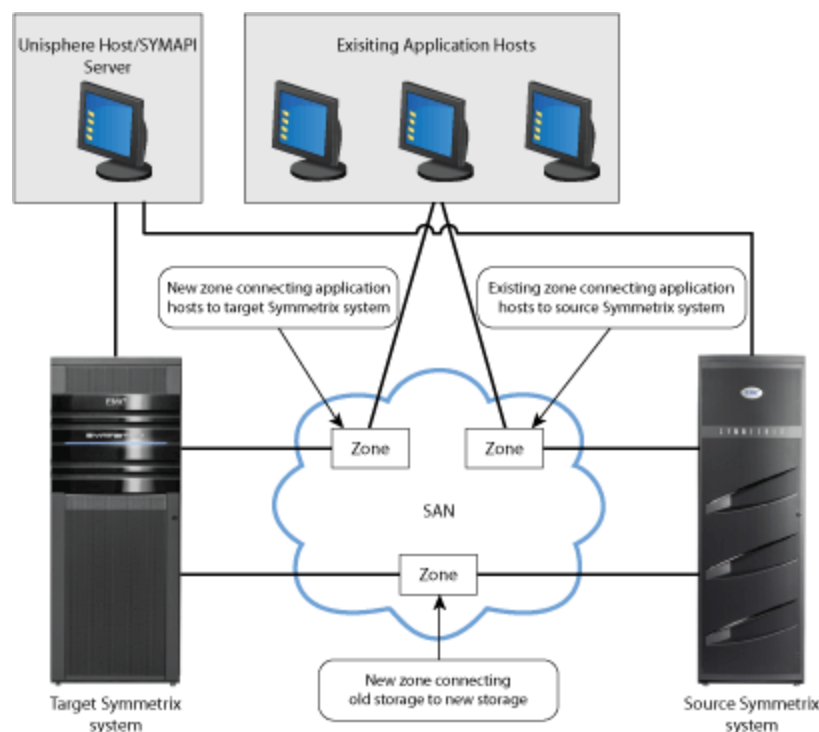
Federated Live Migration

Understanding Federated Live Migration

Federated Live Migration (FLM) allows you to move data from older storage into a new Symmetrix VMAX Series system nondisruptively. The host application cutover to use the new Symmetrix VMAX Series volumes is made transparent by a combination of presenting the VMAX Series volumes as additional paths to the old volumes and managing which paths are active through a multipath IO (MPIO) driver on the host. FLM supports PowerPath as the application MPIO driver, and will support additional MPIO drivers in the future. FLM supports moving the data with Open Replicator SAN-based replication, and will support other underlying technologies in the future. Unlike application host-based PPME, control of the migration and cutover is managed through the Symmetrix VMAX Series system. FLM greatly simplifies migrations requiring no remediation when migrating pre-qualified stacks.

Note For information on supported operating systems, file systems, and logical volume managers, refer to the EMC Federated Live Migration Simple Support Matrix available at <http://powerlink.emc.com>.

An example FLM configuration including the network, storage systems, application hosts, and EMC® Unisphere for VMAX host is shown below:



EMC® Unisphere for VMAX includes a Federated Live Migration wizard to guide you through the process of creating the Federated Live Migration session and the necessary masking view; however, you must set up the required zones for before using the wizard. This includes the zones from the application hosts to the target Symmetrix system and from the source Symmetrix system to the target Symmetrix system. For instructions, refer to

Setting up/Running Federated Live Migration

Before you begin:

Data migrations are often complex operations and require careful planning and execution of predetermined procedures. Failure to identify and perform all steps sequentially or work within supported configurations can result in data unavailability or loss.

Due to requirements of multipath environments, you must correctly configure paths to the new Symmetrix VMAX Series devices to maintain data availability. For details on properly configuring your host environment, refer to EMC Federated Live Migration Technical Overview Technical Notes available at <http://powerlink.emc.com>.


Procedure overview:

There are three primary tasks involved with setting up an running a Federate Live Migration:

- ◆ Readyng the environment
- ◆ Creating the migration session and masking view
- ◆ Migration the data

To ready the environment for a migration:


1. Verify that the application is running on the source Symmetrix system and that the Symmetrix system is zoned to the application and SYMAPI servers, as shown next:



2. Create/modify zoning to support connectivity of the target Symmetrix system to the:

- ◆ Application Server (application access to the target Symmetrix system)
- ◆ SYMAPI server (control of the target Symmetrix system)
- ◆ Source Symmetrix system (application data transfer)

As shown next:



Creating the migration session and masking view

Start the **Federated Live Migration** wizard. The wizard will guide you through the process of setting up a Federated Live Migration session, including defining the storage, port, and initiator groups; and the masking view. For instructions on using the wizard, see [Creating a FLM session](#) on next page.

Once you have finished the wizard, you will have a Federated Live Migration session and a masking view that will provide a data path for the migrated host application. At this point, you should:

1. Verify that the application server has discovered the devices and that the multipathing software has created the paths. In addition, you should also verify that the ORS ceiling parameter is set to 100%, unless it is adversely affecting applications that share the FA port. For more information, see [Setting Open Replicator ceiling](#) on page 419.
2. Continue with the remaining steps in this procedure.

Migrating the data

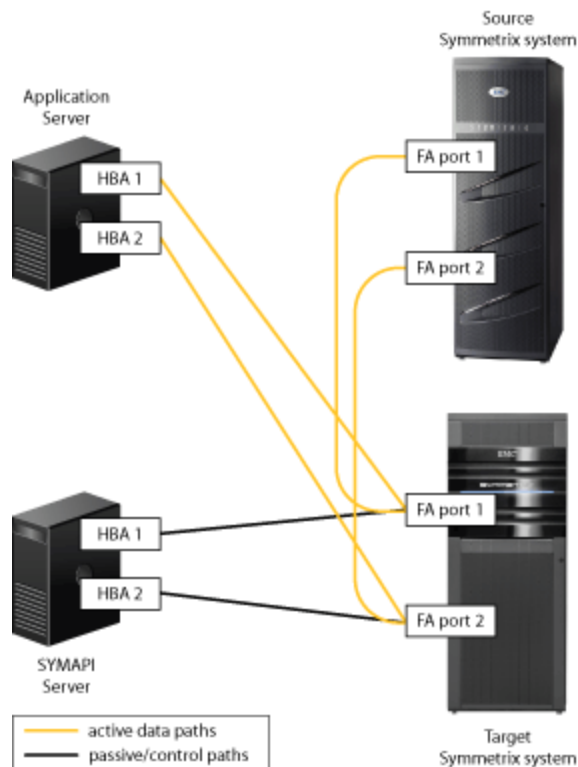
To migrate the data:

1. Activate the Federated Live Migration session.

- Once the migration completes (i.e., all the data has been transferred to the target Symmetrix system), terminate the Open Replicator session, as described in [Terminating Open Replicator sessions](#) on page 419, and remove the zoning and masking of the source Symmetrix ports and volumes. To reset the identities of the volumes used in the migration session, use the Solutions Enabler `symconfigure` command. For instructions, refer to the *EMC Solutions Enabler Symmetrix Migration CLI Product Guide*.

Note The host may run indefinitely with federated identity on the new Symmetrix VMAX Series volumes; however, EMC recommends that the spoofed identity be removed as soon as it is practical to do so following the migration. Leaving the identity spoofing in place long-term has the potential to cause confusion for systems administrators or other users who may not be familiar with the details of FLM and how volume identities are federated. This recommendation is provided only to raise awareness and is not intended as a mandate for unspoofing; Symmetrix VMAX Series volumes may remain federated indefinitely and there is no requirement to unspoof at any time.

The end result is a source Symmetrix system no longer connected to the application or SYMAPI server, as shown below:



Creating a FLM session

There are many environment rules and specifications for running Federated Live Migration (FLM) sessions. Refer to the *Solutions Enabler Symmetrix Symmetrix Data Migration CLI Product Guide* before creating a session.

To create a FLM session:

1. Select **Data Protection > Open Replicator > Open Replicator SAN View** to open the **Open Replicator SAN View**.
2. Filter the remote volumes, as described in [Managing Open Replicator sessions](#) on page 415.
3. Select the volumes and click **Create FLM Session** to open **Create FLM Session** to open the **Federated Live Migration** wizard.
4. Select the **Volumes Source**.
5. *Non-array source only:* Type or select a **Source Name**.
6. Click **Next** and follow one of the options described next.

Auto-match source and target volumes:

- a. Click **Match Selected** to pair your selected source volumes with target volumes.
- b. Click **Match All** to pair all the listed source array volumes with target volumes.

The volume pairs display in the **Selected volume pairs** table.

Note If needed, use the **Remove** button to clear all selected pairs from the table.

Manual selection of source and target volumes:

- a. Select volumes from the **Source array** table and volumes from the **Target array** table and click **Add Pair**.
- b. Select volumes from the **Source array** table and click **Match Selected** to automatically create pairs.

The volume pairs display in the **Selected volume pairs** table.

Note If needed, use the **Remove** button to clear all selected pairs from the table.

7. Click **Next** to open the **Specify Session Options** page.
8. Type a **Session Name** and select the **Host OS**.
9. Type the **HBA type** and **MP type** (these values may be required based on the selection of the Host OS.)
10. *Optional:* Select **Front End Zero Detect**. This feature provides improved performance of pull sessions to thin devices through the detection of incoming zero patterns.
11. Click **Next** to open the **Create Target Storage Group** page.
12. *Optional:* Type a **Storage group name** to change the generated storage group name.
13. Click **Next** to open the **Create Target Port Group** page and follow one of the options described next:
 - ◆ **Create new:** Enter a **Port Group Name**, select the ports to add to the port group.
 - ◆ **Select existing:** This option automatically selects from a list of available ports.

14. Click **Next** to open the **Create Target Initiator Group** page and follow one of the options described next:
 - ◆ **Create new:** Enter an **Initiator Group Name**, and select the initiators/initiator groups to add to the initiator group.
 - ◆ **Select existing:** This option automatically selects from a list of available initiators.

At least one initiator group must be selected in the Initiator/Initiator Group table.

15. Click **Next** to open the **Create Target Masking View** page.
16. *Optional:* Enter a **Masking view name** to change the generated masking view name and click **Finish**.

CHAPTER 7

Performance

This chapter contains the following:

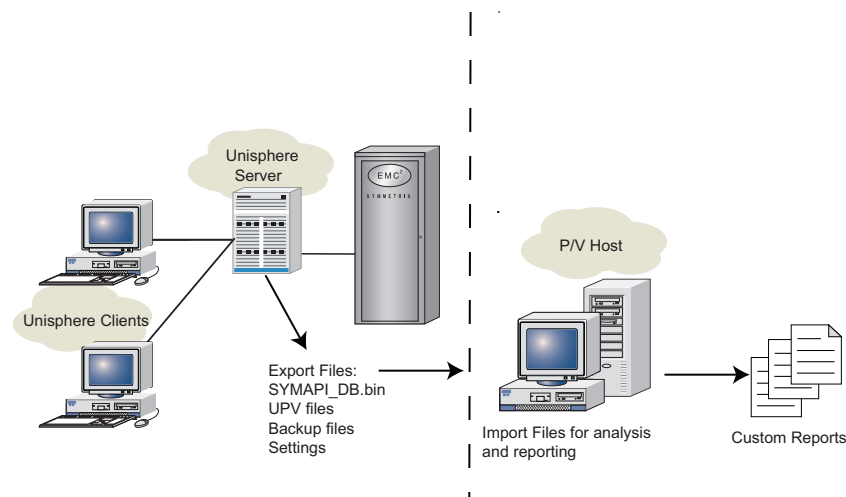
- Using the Performance Viewer 432**
- Monitor view 433**
- Analyze view 444**
- Settings 570**

Using the Performance Viewer

The EMC Unisphere for VMAX 1.6 release includes a Performance Viewer vApp which can be run on an offline host to:

- ◆ Analyze/troubleshoot diagnostic data off site
- ◆ Use historical data for long-range capacity planning
- ◆ Produce reports off site

The Performance Viewer takes the Unisphere data and provides the reporting and diagnostic tools for offline analysis and planning.



To use the Performance Viewer, follow the installation and configuration information in the *EMC Unisphere Performance Viewer Installation Guide*.

The following file types can be exported from Unisphere for VMAX and imported to the Performance Viewer:

- ◆ SYMAPI databases UPV files (diagnostics and real time traces)
- ◆ Backup files (historical data, diagnostics data, and Real Time traces if applicable)
- ◆ Settings (thresholds, reports, user-defined dashboards, and templates)

On the Unisphere side, all files are located in the following directories on the server side:

SYMAPI Database — `SMAS\Backup\SPA\Viewer\SE`

Settings — `SMAS\Backup\SPA\Viewer\Settings`

Monitor view

Managing dashboards

A dashboard is a collection of charts that you define, or that Unisphere for VMAX provides. A user-defined dashboard must have a unique name, up to 32 characters. In addition a user-defined dashboard can be saved to a named folder, or to the default Dashboards folder. Folders can be used to separate systems, applications, or any other object. The folder name displays in the Dashboards tree.

The tree view for the dashboards provides the organization of the dashboard folders. Unisphere for VMAX provides folders for **User Dashboards**, **User Templates**, and **EMC Dashboards**.

Charts and heatmaps that you create can be saved, edited, and deleted. EMC dashboards are created dynamically. You can save the EMC dashboards as a template.

To manage dashboards:

[Creating a dashboard folder \(page 433\)](#)

[Creating a dashboard with charts \(page 434\)](#)

[Creating a dashboard for FAST \(page 435\)](#)

[Creating a heatmap dashboard \(page 434\)](#)

[Copying a dashboard \(page 436\)](#)

[Saving a dashboard as a template \(page 438\)](#)

[Viewing dashboards \(page 438\)](#)

[Deleting a dashboard \(page 437\)](#)

See [Navigating from a heatmap to Analyze](#) on page 442.

To manage EMC dashboards:

[Managing EMC dashboards \(page 439\)](#)

To manage dashboard reports:

See [Running a report from the dashboard](#) on page 442.

See [Scheduling a report from the dashboard](#) on page 442.

Creating a dashboard folder

Before you begin:

You can create a new folder for the dashboards, or create a new folder during the create dashboard process.

To create a dashboard folder:

1. Select **Performance** > **Monitor** to open the **Monitor** view.

2. Click **New Folder**.
3. Enter a **Folder Name** and click **OK**.

The folder displays in the upper pane of the dashboard tree. To create a dashboard that will automatically reside in this folder, select the folder and click **Create**.

Creating a dashboard with charts

Before you begin:

You can create a new folder for the dashboards, or create a new folder during the create dashboard process.

To create a dashboard with charts:

1. Select **Performance** > **Analyze** to open Analyze view.
2. Follow the steps in [Creating charts](#) on page 448.
3. Select **Performance** > **Monitor** to open the **Monitor** view.
4. Click **Create** to open the **Create Dashboard Wizard**.
5. Enter a **Dashboard Name**.
6. Select a **Folder**, or choose **Create a New Folder** and enter a folder name.
7. *Optional:* Set this dashboard as the default. This means that this dashboard displays each time you enter Monitor view. Each user can have only one default dashboard. A Symmetrix system heat map is the default until you specify a default dashboard.
8. Click **Next** to open page 2 of the **Create Dashboard Wizard**.
9. Select **Performance Analyze**, and choose the charts (from step 2) in the Available Charts table.
10. Click **Next** to view a summary of your selections.
11. Click **Finish** to save this dashboard. Click **Back** to make any changes. Click **Cancel** to close the dialog without saving the dashboard.

Creating a heatmap dashboard

Before you begin:

You can create a new folder for the dashboards, or create a new folder during the create dashboard process.

To create a Symmetrix system heatmap dashboard:

1. Select **Performance** > **Monitor** to open the **Monitor** view.
2. Click **Create** to open the **Create Dashboard Wizard**.
3. Enter a **Dashboard Name**.
4. Select a **Folder**, or choose **Create a New Folder** and enter a folder name.

5. *Optional:* Set this dashboard as the default. This means that this heatmap displays each time you enter Monitor view. Each user can have only one default dashboard. A Symmetrix system heat map is the default until you specify a default dashboard.
6. Click **Next** to open page 2 of the **Create Dashboard Wizard**.
7. Select **Predefined Dashboard > Heat Map**, and click **Next** to open page 3 of the **Create Dashboard Wizard**.
8. Enter a dashboard **Title**, select the **Symmetrix ID**, and select the dashboard **Type** (Diagnostic, Historical, Real Time).
9. Set the **Time Range**. Preset values display for each dashboard Type.
Optional: If you select a custom time range, an additional dialog displays. Enter the **Start Time**, **End Time**, and click **OK**.
10. Click **Next** to open page 4 of the **Create Dashboard Wizard**. This page shows a summary of your selections.
11. Click **Finish** to save this dashboard. Click **Back** to make any changes. Click **Cancel** to close the dialog without saving the dashboard.

See ["Navigating in Analyze view" on page 444](#)

Creating a dashboard for FAST

Before you begin:

You can create a new folder for the dashboards, or create a new folder during the create dashboard process.

To create a dashboard for FAST:

1. Select **Performance > Monitor** to open the **Monitor** view.
2. Click **Create** to open the **Create Dashboard Wizard**.
3. Enter a **Dashboard Name**.
4. Select a **Folder**, or choose **Create a New Folder** and enter a folder name.
5. *Optional:* Set this dashboard as the default. This means the FAST dashboard displays each time you enter Monitor view. Each user can have only one default dashboard. A Symmetrix system heat map is the default until a default dashboard is selected.
6. Click **Next** to open page 2 of the **Create Dashboard Wizard**.
7. Select **Predefined Dashboards** and choose one of the FAST options. Dashboard options are **FAST by Storage Group** and **FAST by Tier**.
8. Click **Next** to open page 3 of the **Create Dashboard Wizard**.
9. Select the **Symmetrix ID**, **Storage Group** name, and the **Time Range**. If you select a custom time range, an additional dialog displays to set the Start Time and End Time.
10. Click **Next** to open page 4 of the **Create Dashboard Wizard**. This page shows a summary of your selections.
11. Click **Finish** to save this dashboard. Click **Back** to make any changes. Click **Cancel** to close the dialog without saving the dashboard.

Editing a template dashboard

Before you begin:

You must have a configured dashboard template.

To edit a template dashboard:

1. Select **Performance** > **Monitor** to view the list of User Template dashboards.
2. Select the user template dashboard and click **Edit** to open the Edit Template dialog box.
3. *Optional:* Change the dashboard **Name**.
4. *Optional:* Change the dashboard **Folder** or set the dashboard as the default.
5. *Optional:* Change the value in **Instances**.
6. *Optional:* Select an item from the **Chart List** and click **Edit** to open the Edit Chart dialog box or click **Delete** to delete the chart from the list. Deleting a chart requires a confirmation **OK**.
7. Change the chart **Title**, **Style**, **Type**, or **Format**, **Time Range** and click **OK**.
8. Click **OK**.

Copying a dashboard

You can copy an existing dashboard and change the name and a few characteristics to make a new dashboard.

Before you begin:

You must have an existing dashboard in the **Performance, Monitor** view.

To copy a dashboard:

1. Select the dashboard to copy and click **Copy** to open the Copy Dashboard dialog box.
2. Enter a new **Dashboard Name**.
3. Select a **Folder** for the dashboard location.
4. *User Dashboards only:* Set this dashboard as the default. This means that this dashboard displays each time you enter Monitor view. Each user can have only one default dashboard. A Symmetrix system heat map is the default until you specify a default dashboard.
5. *Template Dashboards only:* Select the **Instances** for the new dashboard.
6. *Optional:* Select a chart from the **Chart List** and use **Delete** or **Edit** to customize the copied dashboard.
7. Click **OK**.

Editing a dashboard

Before you begin:

You must have a configured dashboard.

To edit a dashboard:

1. Select **Performance** > **Monitor** to view the list of dashboards.
2. Select the dashboard and click **Edit** to open the Edit Dashboard dialog box.
3. *Optional:* Change the dashboard **Folder** or set the dashboard as the default.
4. Select an item from the Chart list and click **Edit** to open the Edit Chart dialog box.
5. Change the chart **Title**, **Type**, or **Time Range** and click **OK**.
6. Click **OK**.

Deleting a dashboard

Before you begin:

You can only delete a user-defined existing dashboard. Predefined dashboards cannot be removed.

To delete a dashboard:

1. Select **Performance** > **Monitor** to open Monitor view.
2. Select the dashboard from the top of the dashboard navigation tree and click **Delete**.
3. Click **OK** on the delete confirmation message.

Viewing dashboards

Before you begin:

Read [Managing dashboards](#) on page 433.

To view dashboards:

1. Select **Performance, Monitor** to open Monitor view.

The Monitor view is divided into two sections:

- ◆ **Folders** — The left section of the view displays folders and dashboards.
 - **User Dashboards** — This default folder is the container for user-defined dashboards. You can also create additional folders for user-defined dashboards.
 - **User Templates** — This folder is the container for any dashboard that you want to save. For example, you can configure any of the EMC Dashboards dynamically, and then save it as a template to eliminate the need to configure it again.
 - **EMC Dashboards** — This folder contains predefined dashboards. These dashboards can be modified dynamically, but cannot be deleted. They can be modified and saved as a template.
- ◆ **Dashboard View Block** — The right section of the view displays the selected dashboard. The top of the view block provides dynamic configuration controls, when applicable.

The following controls are available:

- ◆ New Folder — See [Creating a dashboard folder](#) on page 433.
- ◆ Create Dashboard — See [Creating a dashboard with charts](#) on page 434.
- ◆ Save as Template — See [Saving a dashboard as a template](#) below.
- ◆ Run Report Now — See [Running a report from the dashboard](#) on page 442.
- ◆ Schedule Report — See [Scheduling a report from the dashboard](#) on page 442.
- ◆ Edit — See [Editing a dashboard](#) on page 436.
- ◆ Save — See [Saving dashboard changes](#) on the facing page.
- ◆ Copy — See [Copying a dashboard](#) on page 436.
- ◆ Delete — See [Deleting a dashboard](#) on previous page.
- ◆ Navigate to Details View — [Navigating to the Details view](#)

Saving a dashboard as a template

Each of the EMC-defined dashboards can be adjusted and saved as a template for your own dashboard.

To save a dashboard as a template:


1. Select the Symmetrix system, or select All Symmetrix systems.


2. Select **Performance**, and then **Monitor** to open the **Monitor** view.
3. Select one of the **EMC Dashboards**.
4. Set the dashboard characteristics (Type, Time Range) at the top of the dashboard.
5. Click Save as Template to open the **Save Template** dialog box.
6. Enter a dashboard **Name**, select a **Folder**, and select the **Instances** to monitor.
7. *Optional:* If you have charts saved from Analyze view, click **Add** to add them to the dashboard.
8. Click **OK**.

Saving dashboard changes

The **Save** option in the Monitor view allows you to save changes that you make dynamically to a dashboard chart using the chart controls.

To save dashboard changes:

1. Open Unisphere for VMAX to **Performance > Monitor**.
2. Select a dashboard.
3. Using the individual chart icons, make your changes to the chart properties  (**Chart Style, Data Format, Display Thresholds**), delete the chart (x), or change the dashboard columns.



4. Click **Save**.

Managing EMC dashboards

To view and manage EMC dashboards:

1. Select **Performance > Monitor** to open the Monitor view.
2. Select a dashboard from the **EMC Dashboards** folder.
3. Make your selections from the fields across the top of the dashboard.

The charts update dynamically based on your selections.

Note To keep a predefined dashboard, use the **Create** (Dashboard wizard) and select **Predefined Dashboards** as the source. All the EMC Dashboards are available for selection in the wizard.

Using the [FAST dashboards](#) on next page provides more details about the EMC dashboards for monitoring FAST data movement between the tiers.

Using the FAST dashboards

When you are monitoring the FAST activity, the best place to start is with the **FAST by Policy** dashboard, then with the **FAST by Storage Group** dashboard. The **FAST by Tier** dashboard provides another view of the system at the tier level.

Note All the EMC FAST dashboards in **Performance > Monitor** are for FAST VP.

- ◆ The **FAST by Policy** dashboard shows a high-level roll-up of the storage groups and tiers. You can use this information for capacity planning and trending.
[Using the FAST by Policy dashboard](#) (page 441).
- ◆ The **FAST by Storage Group** dashboard narrows down the view to the performance of a specific storage group.
[Using the FAST by Storage Group dashboard](#) (page 441).
- ◆ The **FAST by Tier** dashboard shows the specified tier performance - which can include many storage groups per tier.
[Using the FAST by Tier dashboard](#) (page 440).

Using the FAST by Tier dashboard

The **FAST by Tier** dashboard a different view of the FAST management. In this dashboard the main entity is the virtual pool tier. The charts break down the activity and allocations within the tier by the storage groups that are using the selected tier.

You select the **Symmetrix ID**, **Virtual Pool Tier**, view **Type (Diagnostic or Historical)**, and **Time Range** for the dashboard.

The following are descriptions of the **FAST by Tier** dashboard charts:

- ◆ **Tier Ingress/Egress** — Shows the number of extents entering and leaving the virtual pools that reside on the tier. It shows that the FAST controller is moving extents around. There is no good or bad value for this chart. It is informational only.
- ◆ **Allocation by Storage Group** — Shows the capacity of each storage group within the tier. A tier can have hundreds of associated storage groups; the chart only shows the first several.

Hint Change to the **FAST by Storage Group** dashboard if you want to focus on a specific storage group.

- ◆ **BE MBs Transferred by SG** — Shows the throughput rate to/from the tier by each of the storage groups.
- ◆ **BE Requests by SG** — Shows the request rate to/from the tier by each of the storage groups.
- ◆ **IO Density by SG** — Shows how efficiently the FAST controller is managing the IO of each of the storage groups.
- ◆ **Read and Write Response Time** — Shows the time it takes the disk to satisfy IO requests.

Using the FAST by Storage Group dashboard

The **FAST by Storage Group** dashboard shows you how FAST is managing your application data. FAST performance is only measured by the IO activity on the back-end disks. Host IO that is satisfied from cache, does not factor into FAST performance metrics.

You select the **Symmetrix ID**, **Storage Group**, view **Type (Diagnostic or Historical)**, and **Time Range** for the dashboard.

These are the chart descriptions for the **FAST by Storage Group** dashboard:

- ◆ **Allocation by tier** — Shows how much of the storage group is allocated to each tier. You want to see most of the data on SATA, less than that on Fiber, and less than that on EFD.
- ◆ **IO Density by tier** — Shows how many IOs are done per GB of disk. FAST is successful when the EFD value is the highest, the SATA value is the lowest, and the Fiber value is in the middle.
- ◆ **BE Requests/sec by tier** — Shows the number of requests made each second to each tier. You want to see most of the requests either on EFD or trending towards EFD.
- ◆ **BE MBs/sec by tier** — Shows the throughput made each second to each tier. You want to see most of the throughput either on EFD , or trending towards EFD.
- ◆ **Read and Write Response Times** — Shows the time it takes to satisfy IO requests. This is the proof of success chart for FAST. You want to see your response times improving over time.
- ◆ **Host IOs/sec and BE Requests/sec** — This chart compares the host IOs that are satisfied from cache with the back-end requests that FAST is handling. Workloads that serve all IOs from cache are usually not good candidates for FAST management; applications that have a low hit ratio benefits most from the FAST management.

Using the FAST by Policy dashboard

A FAST policy defines a limit for each tier in the policy. This limit determines how much capacity from a storage group associated with the policy is allowed to reside on each tier.

The **FAST by Policy** dashboard shows you how FAST is managing your data associated with a specific policy. You select the **Symmetrix ID**, **FAST Policy**, view **Type (Diagnostic or Historical)**, and the **Time Range** for the dashboard.

The following are descriptions of the **FAST by Policy** dashboard charts:

- ◆ **FAST Policy SG Capacity** — Shows the total capacity and allocated capacity for the sum of all the storage groups associated with the selected policy. This is a logical view that shows you how much you are using and how much you can expect to grow.
- ◆ **FAST Policy Tier Capacity** — Shows the total capacity (what you have) and allocated capacity (what you are currently using) for the sum of all the tiers associated with the selected policy. This is a physical show you how much room you have to grow.
- ◆ **BE Requests** — Shows the sum of all back-end requests done by all storage groups associated with the selected policy.
- ◆ **BE Throughput** — Shows the sum of all back-end throughput done by all storage groups associated with the selected policy.

- ◆ **FAST Internal Activity** — Shows the total number of extents moving in and out each second from the virtual pools governed by the selected policy.
- ◆ **IO Density** — Shows how many IOs are done per GB of disk. FAST is successful when the EFD value is the highest, the SATA value is the lowest, and the Fiber value is in the middle.

Scheduling a report from the dashboard

The performance dashboard reports are saved in .pdf format. When you schedule a report from the Monitor dashboard view, it will run continuously on the time/days selected.

To schedule a report from the dashboard:

1. Select **Performance**, and then **Monitor** to open the **Monitor** view.
2. Select or create a system dashboard.
3. Click **Schedule Report** to open the Schedule Report dialog box.
4. Enter a report **Name**.
5. *Optional:* Enter the report description.
6. Select the **Execution Time** and the **Days to run**.
7. Click **OK**.

A confirmation message displays. Any changes to this scheduled report are made in **Settings > Reports**.

Scheduled reports save to this default location:

```
SMAS\jboss\standalone\data\spa\querydata\
```

Running a report from the dashboard

The performance dashboard reports are saved in .pdf format. Dashboard reports show the charts as they appear in the dashboard view. To create reports for other data formats (xml, csv, btp), and to manage any scheduled reports, go to **Settings > Reports**.

To run a report from the dashboard:

1. Select **Performance**, and then **Monitor** to open the **Monitor** view.
2. Select or create a system dashboard.
3. Click **Run Report Now**.
4. Click **OK** to save the query results.
5. Type a file name in the **Save File** dialog box.

The report is saved to the specified location.

Note Reports that are "Run Now" do not appear in the **Reports** list.

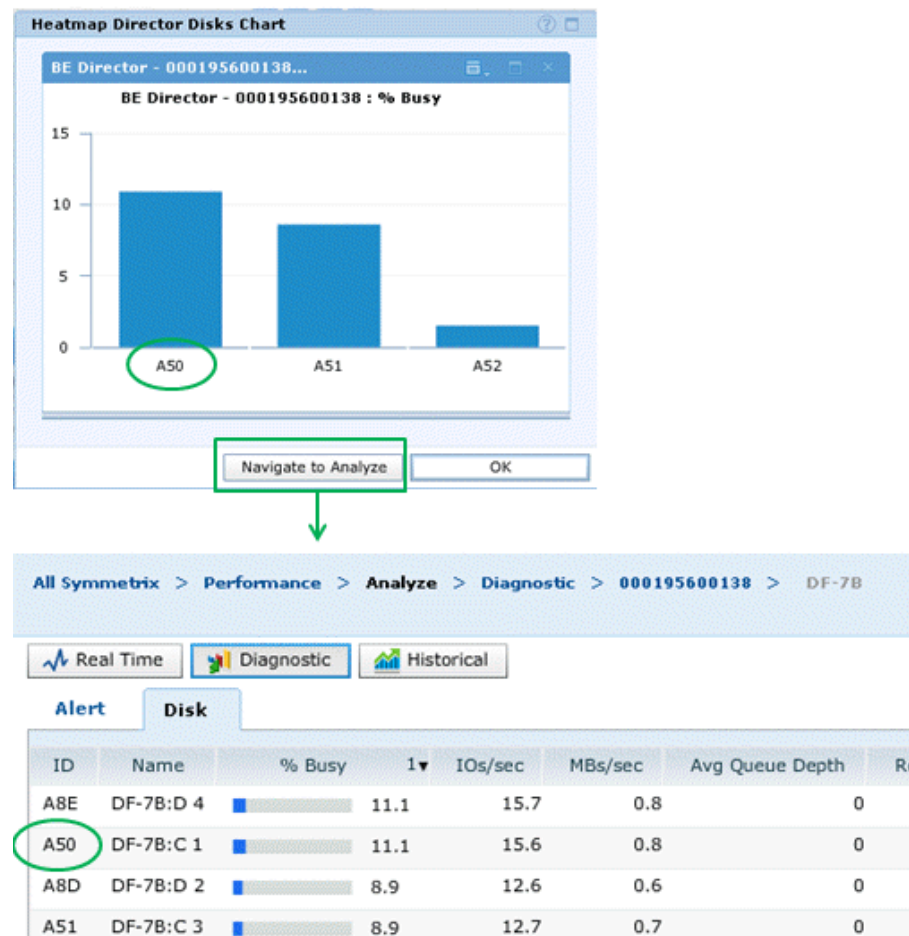
Navigating from a heatmap to Analyze

You can navigate from an EMC Symmetrix heatmap object to Analyze Diagnostic view.

To navigate to Analyze from a heatmap:

1. Select the Symmetrix system.
2. Select **Performance** > **Monitor** to open the **Monitor** view.
3. Select **EMC Dashboards** > **Heatmap**.
4. *Optional:* Select the **Type** (**Real Time**, **Diagnostic**, or **Historical**) and the **Time Range** from the drop-down options above the heatmap.
5. Double-click on any object in the heatmap to open its corresponding chart.
6. Click **Navigate to Analyze**.

In this example, the back-end disk director DF-7B, A50.



The Analyze, Diagnostic view opens to the selected object's location.

Analyze view

Navigating in Analyze view

In Analyze view, you can view Symmetrix system data for various collection ranges.

- ◆ **Real Time** —The previous 1 hour. Real Time view collects data between 2 and 5 seconds for a limited group of metrics. The data is available for the previous hour.
- ◆ **Diagnostic** — The previous 4 hours. Diagnostic view collects data every 5 minutes for root cause analysis. The data is available for the previous 7 days.
- ◆ **Historical** — The previous 24 hours. Historical view collects data in 15 minute intervals for trending and planning. The data is available for the previous year.

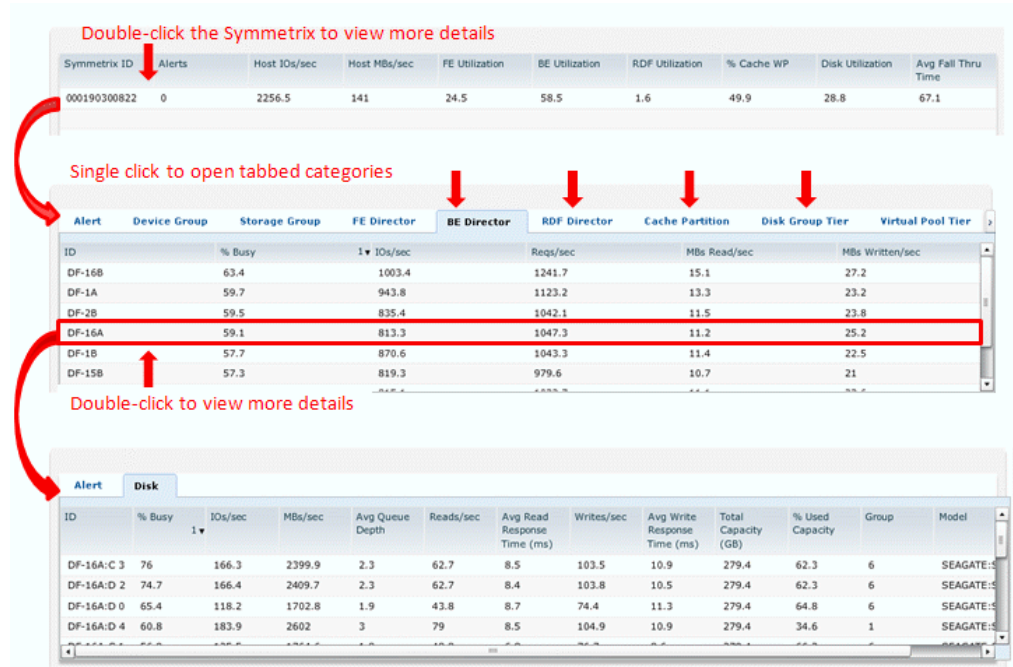
In Analyze view you can do the following:

- Monitoring data (page 444).
- Creating charts (page 448).
- Changing the time range (page 449).
- Creating a dashboard from Explore (page 450).

Monitoring data

The Analyze view provides three views into your storage data: Real Time, Diagnostic, Historical.

Licensed and registered Symmetrix systems display in a table format. To select an item, use one-click; to look at more details, double-click.



About charts

When you create a chart there are some options available for customization.

The **Properties and Settings** icon provides options for the following:

- ◆ Chart Style
- ◆ Data Format
- ◆ Display Threshold
- ◆ Add a metric to a chart
- ◆ Adjust the chart view by columns

Chart Style

The available chart styles are: line, bar, stacked bar, area, pie chart and table, depending on the specified metrics.




Figure 1: Chart styles

Data Format

By default, Performance displays each metric value's Average in the table rows and in the charts. You can change the chart format by selecting one or both of the following:

- ◆ **Average** — Shows the average metric value for the session (hour). This is the default.
- ◆ **Maximum** — Shows the maximum metric value—the high water mark—for the session (hour).

Selecting both formats with a multiple-metrics chart can make a chart less readable.




Figure 2: Data format

The format selected in the charts updates the format in the table cells.

When both formats are selected the table cells display as **Maximum/Average**. For example Host MBs/sec would display 193.9/141.1.

Symmetrix ID	Alerts	Host I/Os/sec	Host MBs/sec	FE Utilization	BE Utilization	RDF Utilization	% Cache WP	Disk Utilization	Avg Fall Thru Time
000190300822	N/A/0	2506.7/2257	193.9/141.1	94.7/24.5	71.2/58.4	21.6/1.6	50.1/49.9	98.5/28.8	69/67

Real Time view only

Latest — Displays the latest data in addition to the data averages, as **Latest / Average**. When Maximum is also selected the values display as **Latest / Maximum / Average**.

Capture Trace — Captures the Real Time performance data for the previous hour.

Display Threshold

If you have set thresholds for a metric, you can display the threshold value in the chart.





Figure 3: Display threshold

Add a metric to a chart

Any metric can be added to an existing chart by selecting the metric and selecting **Add to Chart** from the chart properties and settings icon.

Adjust the chart view by columns

Use the column icon drop-down to select the number of columns to display.



Using Zoom icons

Each Performance chart includes an icon (left) to quickly tell you from which view the chart originates.

In addition, there are three zoom capabilities (right). When you use the zoom function, the start and end times do not change.



Creating charts

The lower portion of the Analyze view provides the capability to create user-definable charts. A list of metrics dynamically changes based on your selections in the table. The metrics list can display:

- ◆ **Key Performance Indicators** — Lists the key performance indicators that correspond to the column headings in the current table.
- ◆ **All** — Lists all available metrics for the selected object in the table.

You can customize the metrics list. See [Metrics](#) on page 574.

Before you begin:

Verify that you have the correct time range for the charts and graphs. See [Metrics](#) on page 574.

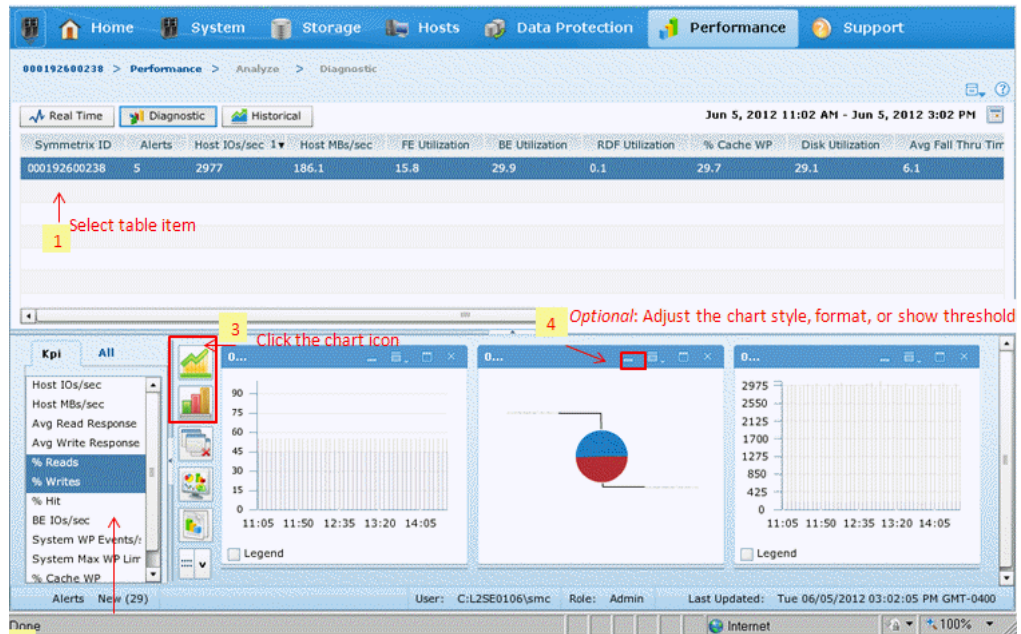
To create charts:

1. Select the system or component from the top table. Multi-select is allowed in the table.
2. Select the metric(s) that you want to use from the **Key Performance Indicator** list or the **All** list (bottom-left). You can select multiple metrics, however, the more you select, the less readable the chart will be.
3. Click the create chart icon. The top icon creates a time series chart which displays the values charted over time. The second icon creates a snapshot chart which displays a single average value over time.

Each chart has icons for properties (style, format, thresholds), maximization, delete, and refresh.

4. *Optional:* Modify the chart style, format, or add the threshold on the chart. See [Viewing thresholds on charts](#) on page 576.

The available chart styles are: line, bar, area, and stacked bar.



Changing the time range

In the **Analyze** view, there are the following default data collection time ranges:

- ◆ Real Time — The previous 1 hour.
- ◆ Diagnostic — The previous 4 hours.
- ◆ Historical — The previous 24 hours.

You can customize the time range for each data collection view.

To change the Real Time time range:

1. Click the calendar icon to the right of the time display to open the Time Selection dialog box.
2. Make one of the following selections:
 - a. Select **Now** to view the last hour of data collected.
 - b. Select **Trace** and choose a saved trace file from the list.
3. Click **OK**.

To change the Diagnostic time range:

1. Click the calendar icon to open the **Time Selection** dialog box.
2. Make one of the following selections:
 - a. Select **Now** and choose the last 1, 2, 4, 8, or 24 hours.
 - b. Select **By Day** and set the **Start Time** and **End Time**.
 - c. Select **Fixed** and click the **Start Time** calendar icon to select a start date.
Adjust the time by one hour increments.
Click the calendar to the right of the **End Time** calendar icon and select an end date.

Adjust the time by one hour increments.

3. Click **OK**.

To change the Historical time range:

1. Click the calendar icon to the right of the time display to open the Time Selection dialog box.
2. Make one of the following selections:
 - a. Select **Now** and choose the last 24 hours, 1 week, 1 month, 6 months, or 12 months.
 - b. Select **Fixed** and click the **Start Time** calendar icon to select a start date.

Adjust the time by one hour increments.

3. Click **OK**.


Creating a dashboard from Explore

You can automatically create a dashboard from Analyze view while you are creating charts for your optimal data concerns.

Before you begin:


You must have access to a visible, registered Symmetrix system.

To create a dashboard from Explore:

1. Select component(s) from the table in Analyze view, and [create your charts](#).
2. Click the Create a dashboard icon.  The Create Dashboard dialog box opens.
3. Enter a **Dashboard Name**.
4. Select a **Folder** for the dashboard.
5. *Optional:* Set this dashboard as the default. This means that this dashboard displays each time you enter Monitor view. Each user can have only one default dashboard. A Symmetrix system heat map is the default until you specify a default dashboard.
6. Click **OK** to save the dashboard.

Creating a template dashboard





To create a dashboard template from Explore:


1. Select component(s) from the table in Analyze view, and [create your charts](#).
2. Click the Create template dashboard icon.  The **Create Template Dashboard** dialog box opens.
3. Enter a **Dashboard Name**.
4. Select a **Folder** for the dashboard.
5. *Optional:* Select the **Instances** to monitor. Not all components have this option.
6. Click **OK** to save the dashboard.

The dashboard saves to the specified folder in **Monitor** view.

Using the icons in Analyze view

The following icons are found in the **Performance, Analyze** view. Each icon also provides a description when you mouse over it.

Icon	Description
	<p>Draws a time series chart using the selected elements. Time series charts display the values charted over time.</p> <p>Draws a snapshot chart using the selected elements. A snapshot chart displays a single average value over time.</p> <p>Deletes all displaying charts.</p> <p>Opens the Create Dashboard dialog box containing the charts in the display.</p> <p>Opens the Create Template Dashboard dialog box containing the charts in the display.</p> <p>Sets the number of columns for the chart view.</p>
	<p>Changes the table layout. Possible values in Diagnostic are:</p> <p>Average — Shows the average metric value for the session (hour). This is the default.</p> <p>Maximum — Shows the maximum metric value—the high water mark—for the session (hour).</p> <p>When you select Maximum, the values display Maximum / Average.</p> <p>Possible values in Real Time are Average, Maximum and the following:</p> <p>Latest — Displays the latest data in addition to the data averages, as Latest / Average. When Maximum is also selected the values display as Latest / Maximum / Average.</p> <p>Capture Trace — Captures the Real Time performance data for the previous hour.</p>
	<p>Opens the Time Selection dialog box. Use this to change the time range in Diagnostic and Historical views. In Real Time view, use this dialog to change the view to a previously captured trace, and return the time range to Now.</p>
	<p>The chart icons from left to right are:</p> <p>Minimize — Minimize the chart.</p>

Icon	Description
	<p>Chart Properties — Opens the following:</p> <p>Chart Style — Styles vary by the selected metric. Possible chart styles are: Line, Bar, Stacked Bar, Area, Table, Pie Chart.</p> <p>Data Format — Possible values are Average and Maximum.</p> <p>Display Thresholds — Adds a line to the chart to show the user-defined thresholds. If a threshold is not defined, this option is dimmed.</p> <p>Add to Chart — Adds the selected metric to the current chart.</p> <p>Trend Line — Adds a trend line to the current chart.</p> <p>Maximize — Maximize the chart.</p> <p>Remove — Removes the chart.</p>
<p>Forecast Points: <input type="text" value="0"/> </p>	<p>This option displays in the chart after you add a Trend Line. It allows you to set the number of data points to use for the trend line.</p>

Analyze Real Time

Symmetrix system Real Time view

Real Time view provides the following tabbed information for the selected Symmetrix system:

Table view:

- ◆ **FE Director** — The front-end director IDs for the selected Symmetrix ID. Select (double-click) a specific front-end director ID table row to change the information display to that director.
- ◆ **BE Director** — The back-end director IDs for the selected Symmetrix ID. Select (double-click) a specific back-end director ID table row to change the information display to that director. From the specific director, you can navigate to the hyper level.
- ◆ **RDF Director** — The RDF directors for the selected Symmetrix ID. Select (double-click) a specific RDF director ID table row to change the information display to that director. From the specific director, you can navigate to the hyper level.

Charts:

[Symmetrix system Real Time metrics](#) (page 548)

Front-end director Real Time view

Real Time view provides the following information about the front-end directors:

Tableview:

- ◆ **ID** — The ID of the director.

- ◆ **% Busy** — The percent of time the directory is busy.
- ◆ **Host IOs/sec** — The IOs per second for each front end director.
- ◆ **Host MBs/sec** — The MBs per second for each front end director.
- ◆ **Reqs/sec** — The number of requests per second for each front end director.
- ◆ **System WP Events/sec** — The number of write pending events per second for the system.
- ◆ **Device WP Events/sec** — The number of write pending events per second for the volume.

Charts:

See [Front-end director Real Time metrics](#) on page 521.

Back-end director Real Time view

Real Time view provides the following information about back-end directors (DA and DX):

Table view:

- ◆ **ID** — The ID of the director.
- ◆ **% Busy** — shows the percent of time the directory is busy.
- ◆ **IOs/sec** — The number of IOs to and from the disks that are mapped to this back-end director.
- ◆ **Reqs/sec** — The requests per second for each back-end director.
- ◆ **Reads Reqs/sec** — The number of read requests to and from the disks that are mapped to this back-end director.
- ◆ **Writes Reqs/sec** — The number of write requests to and from the disks that are mapped to this back-end director.

Charts:

See [Back-end director Real Time metrics](#) on page 491.

RDF director Real Time view

Real Time view provides the following information for RDF directors:

Table view:

- ◆ **ID** — The ID of the director.
- ◆ **% Busy** — The percent of time the directory is busy.
- ◆ **IOs/sec** — The number of IOs to and from the disks that are mapped to this RDF director.
- ◆ **Tracks Sent/sec** — The number of tracks sent to each disk from the RDF director.
- ◆ **Tracks Received/sec** — The number tracks received from each disk mapped to the RDF director.

Charts:

See [RDF director Real Time metrics](#) on page 528.

Viewing a Real Time trace

You can view saved traces of Real Time data to troubleshoot or diagnose performance issues.

Before you begin:

One or more saved data traces.

To view a Real Time trace:

1. Select the Symmetrix system.
2. Select **Performance, Analyze** to open Analyze view.
3. Select **Real Time**.
4. Click the Calendar icon in the Time Range to open the Time Selector dialog box.
5. Select **Trace** and choose a saved trace file from the list.
6. Click **OK**.

When the trace displays in the table, you can create charts for monitoring that time range.

Analyze Diagnostic

Symmetrix systems Diagnostic view

Diagnostic view provides the following information for the list of Symmetrix systems:

Table view:

- ◆ **Symmetrix ID** — The Symmetrix ID for each supported system in the domain.
- ◆ **Alert** — The number and severity of alerts associated with each Symmetrix system. The alerts provide access to the fast lane. After selecting the Symmetrix ID, double-click an alert ID to automatically drill down to the alert issue.
- ◆ **Host IOs/sec** — The number of host IO operations performed each second by all Symmetrix volumes, including writes and random and sequential reads.
- ◆ **Host MBs/sec** — The number of host MBs written and read by all of the Symmetrix volumes each second.
- ◆ **FE Utilization** — The calculated percent of time the front-end directors are busy.
- ◆ **BE Utilization** — The calculated percent of time the back-end directors are busy.
- ◆ **RDF Utilization** — The calculated percent of time the RDF directors are busy.

If the utilization value is not available, the number indicates the MBs per second activity between the local and remote Symmetrix systems, if both systems are in the managed domain. If the remote system is not in the domain, the value for each of the local R1 links displays without a remote identifier.
- ◆ **% Cache WP** — The percent of cache that is busy with write pendings.

- ◆ **Disk Utilization** — The average utilization across the disk groups. The utilization (time busy) values are 0% to 100%.
- ◆ **Average Fall Through** — The average time it takes a cache slot in LRU0 to be freed up. It is the average time from the first use of the contents to its reuse by another address.

Charts:

See [Metric Tables](#) on page 483.

Single Symmetrix system Diagnostic view

Diagnostic view provides the following information for a single Symmetrix system:

Table view:

- ◆ **Alert** — The ID and severity of each alert for the Symmetrix ID. The alerts provide access to the fast lane. Double-click an alert ID to automatically drill down to display the alert issue.
- ◆ **Device Group** — The volume groups that belong to the selected Symmetrix ID, and displays key performance indicators for each volume group in the table and the dashboard charts. Select (double-click) a specific volume group table row to change the information display to that volume group.
- ◆ **Storage Group** — The storage groups configured for the selected Symmetrix ID, and displays key performance indicators for each storage group, both in the table and the dashboard charts. Select (double-click) a specific storage group table row to change the information display to that storage group.
- ◆ **FE Director** — The front end director IDs for the selected Symmetrix ID, and the key performance indicators for each director, both in the table and the dashboard charts. Select (double-click) a specific front end director ID table row to change the information display to that director.
- ◆ **BE Director (DA)**— The back end director IDs for the selected Symmetrix system, and the key performance indicators for each back end director, both in the table and the dashboard charts. Select (double-click) a specific back end director ID table row to change the information display to that director. From the specific director, you can drill-down to the hyper level.
- ◆ **BE Director (DX)**— The back end director IDs for the external disk. From the specific director, you can drill-down to the external disk.
- ◆ **RDF Director** — The RDF directors for the selected Symmetrix system, and the key performance indicators for each RDF director, both in the table and the dashboard charts. Select (double-click) a specific RDF director ID table row to change the information display to that director. From the specific director, you can drill-down to the hyper level.
- ◆ **Cache Partition** — The default and configured cache partitions for the selected Symmetrix ID, and displays key performance indicators for each cache partition, both in the table and the dashboard charts. Select (double-click) a specific cache partition table row to change the information display to that cache partition.
- ◆ **Disk Group Tier** — The disk group tiers configured for the selected Symmetrix ID, and displays key performance indicators for each tier, both in the table and the dashboard charts. Select (double-click) a specific disk group tier table row to change the information display to that tier.

- ◆ **Virtual Pool Tier** — The virtual pool tiers configured for the selected Symmetrix ID, and displays key performance indicators for each tier, both in the table and the dashboard charts. Select (double-click) a specific virtual pool tier table row to change the information display to that tier.
- ◆ **FAST VP Policy** — The storage group capacities and allocations for FAST VP, and displays key performance indicators for each policy.
- ◆ **Snap Pool** — The snap pools configured for the selected Symmetrix system, and the key performance indicators for each snap pool, both in the table and the dashboard charts. Select (double-click) a specific snap pool ID table row to refresh the information display to show only that snap pool. From the snap pool, you can drill down to the related SAVE volumes, disks, and hypers.
- ◆ **TP Pool** — The thin pools configured for the selected Symmetrix system, and the key performance indicators for each thin pool, both in the table and the dashboard charts. Select (double-click) a specific thin pool ID table row to refresh the information display to show only that thin pool. From the thin pool, you can drill down to the related thin and DATA volumes, disks, and hypers.
- ◆ **DSE Pool** — The DSE pools configured for the selected Symmetrix system, and the key performance indicators for each DSE pool, both in the table and the dashboard charts. Select (double-click) a specific DSE pool ID table row to refresh the information display to show only that DSE pool. From the DSE pool, you can drill down to the related SAVE volumes, disks, and hypers.
- ◆ **Disk Group** — The disk groups configured for the selected Symmetrix system, and the key performance indicators for each disk group, both in the table and the dashboard charts. Select (double-click) a specific disk group ID table row to refresh the information display to show only that disk group. From the disk group, you can drill down to the related disks and hypers.
- ◆ **External Disk Group** — The external disk groups configured for the selected Symmetrix ID, and displays key performance indicators for each external disk group, both in the table and the dashboard charts. Select (double-click) a specific external disk group to view its related alerts and external disks.
- ◆ **Host IO Limit by SG** — The IO and bandwidth limits that were set for the storage group.
- ◆ **Host IO Limit by FE** — The IO and bandwidth limits that were set for the front end port.
- ◆ **Event** — The recent events for the selected Symmetrix system. This list is display-only; information about events does not display in the dashboard.

Charts:

See [Symmetrix system metrics](#) on page 544.

Alerts Diagnostic view

Diagnostic view provides the following information about alerts:

- ◆ **Created** — The time the alert was created.
- ◆ **Severity** — The severity of the alert.
 - (1) Fatal
 - (2) Critical

- (3) Warning
- (4) Information
- (5) Normal

- ◆ **Category** — The object related to the alert, such as volume group, disk group, etc.
- ◆ **Instance** — The specific category of the alert. For example, if the category listed disk groups, the instance would name the specific disk.
- ◆ **Metric** — The metric that generated the alert.
- ◆ **Value** — The specific value of the metric that generated the alert.
- ◆ **Count** — The number of times the metric exceeded the threshold.
- ◆ **Acknowledge** — Indicates whether the alert has been acknowledged.
- ◆ **Type** — Indicates whether the alert is static or dynamic.
- ◆ **Message** — Describes the cause of the alert.

Device group Diagnostic view

Diagnostic view provides the following information about device groups:

Table view:

- ◆ **ID** — The name assigned to the device group.
- ◆ **Host IOs/sec** — The number of host operations performed each second by the group.
- ◆ **Read Response Time (ms)** — The average time it took the Symmetrix system to serve one read IO for this group.
- ◆ **Write Response Time (ms)** — The average time that it took the Symmetrix system to serve one write IO for this group.
- ◆ **% Hit** — The percent of IO operations, performed by the group, that were immediately satisfied by cache.
- ◆ **% Writes** — The percent of IO operations that were writes.
- ◆ **% Read Miss** — The percent of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
- ◆ **Capacity (GB)** — The capacity of the device group in GBs.
- ◆ **Members** — The number of volumes that comprise this device group.

Charts:

See [Group metrics](#) on page 502.

Metavolumes Diagnostic view

Diagnostic view provides the following information about metavolumes:

Table view:

- ◆ **ID** — The Symmetrix volume number.
- ◆ **Host IOs/sec** — The number of host read IO and write IO operations performed each second by the Symmetrix volume.
- ◆ **Avg Read Time (ms)** — The average time it takes the volume to perform the Reads in milliseconds.
- ◆ **Avg Write Time (ms)** — The average time it takes the volume to perform the Writes in milliseconds.
- ◆ **% Hit** — The percentage of IO operations, performed by the Symmetrix volume, that were immediately satisfied by cache.
- ◆ **% Write** — The percent of IO operations that were writes.
- ◆ **% Read Miss** — The percent of read miss operations performed each second by the volume. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.

Charts:

See [Volume metrics](#) on page 557.

Storage group by tier Diagnostic view

Diagnostic view provides the following information about storage groups by tiers:

Table view:

- ◆ **ID** — The storage group tier ID.
- ◆ **BE Read Reqs/sec** — The number of read requests each second performed by the disk director to cache.
- ◆ **Destage Writes/sec** — The number of writes per second that were destaged to disk.
- ◆ **IOs/sec** — The number of IOs per second for data transfer.
- ◆ **BE MBs Read/sec** — The number of host read IOs performed by the disk group per second.
- ◆ **Destage Write MBs/sec** — The size (MBs) of writes per second that were destaged to disk.
- ◆ **MBs/sec** — The total MBs read and written per second.
- ◆ **Capacity** — The allocated capacity.

For example, if SG1 is 100 GB on Tier1 and 50 GB on Tier2, then the SG capacity is 100 GB for Tier1 and 50 GB for Tier2.

- ◆ **IO Density** — The number of BE requests per GB of disk.
(BE Reads + destaged writes) / capacity

Charts:

See [Disk group storage tier metrics](#) on page 514.

See [Virtual pool tier metrics](#) on page 565.

Tier by storage group Diagnostic view

Diagnostic view provides the following information about tiers by storage groups:

Table view:

- ◆ **ID** — The storage group name.
- ◆ **BE Read Reqs/sec** — The number of read requests each second performed by the disk directors to cache.
- ◆ **BE Write Reqs/sec** — The number of write requests each second performed by the disk directors to cache.
- ◆ **BE Reqs/sec** — The number of requests each second performed by the disk directors to cache.
- ◆ **BE MBs Read/sec** — The number of host read IOs performed by the disk group per second.
- ◆ **BE MBs Written/sec** — The number of host write IOs performed by the disk group per second.
- ◆ **BE MBs Transferred/sec** — The number of MBs performed by the disk group per second.
- ◆ **Allocated Capacity** — The number of GBs of the storage group that are allocated to the FAST policy.
- ◆ **IO Density** — The number BE requests per GB of disk. (BE Reads + destaged writes) / capacity

Charts:

See [Tier by Storage Group metrics](#) on page 551.

FAST VP Policy Diagnostic view

Diagnostic view provides the following information about FAST policies:

Table view:

- ◆ **ID** — The policy ID.
- ◆ **Allocated SG Capacity (GB)** — The number of GBs of the storage group that are allocated to the FAST policy.
- ◆ **Total SG Capacity (GB)** — The total capacity of the storage group.
- ◆ **Allocated Pool Capacity (GB)** — The number of GBs of the virtual pool that are allocated to the FAST policy.
- ◆ **Total Pool Capacity (GB)** — The total capacity of the virtual pool.
- ◆ **BE Reqs/sec** — The number of data transfers between cache and the director.
- ◆ **BE MBs Transferred/sec** — The number of MBs transferred each second between cache and the director.

Charts:

See [FAST policy metrics](#) on page 516.

External disk Diagnostic view

Table view:

- ◆ **ID** — The numeric identifier of the external disk.
- ◆ **Name** — The external disk name.
- ◆ **IOs/sec** — The number of host IO commands per second for data transfer.
- ◆ **MBs/sec** — The throughput number (MBs) per second from the host.
- ◆ **Reads/sec** — The number of read IOs performed by the disk per second.
- ◆ **Avg. Read Response Time (ms)** — The average time it took the disk to serve one read command.
- ◆ **Writes/sec** — The number of write IOs performed by the disk per second.
- ◆ **Avg. Write Response Time (ms)** — The average time it took the disk to serve one write command.
- ◆ **Total Capacity (GB)** — The total capacity of the external disk.
- ◆ **Used Capacity (%)** — The used capacity of the external disk.
- ◆ **Members** — The number of members for the external disk.

Charts:

See [External disk metrics](#) on page 499.

External disk group Diagnostic view

Diagnostic view provides the following information about Symmetrix external disk groups:

Table view:

- ◆ **ID** — The identifier assigned to this external disk group.
- ◆ **Name** — The name of the external disk group.
- ◆ **Reads/sec** — The average number of host reads performed per second.
- ◆ **MBs Read/sec** — The throughput per second for the disks that comprise this disk group.
- ◆ **Avg Read Response Time** — The average time it takes the Symmetrix system to perform read operations.
- ◆ **Writes/sec** — The average number of host writes performed per second.
- ◆ **MBs Written/sec** — The number of MBs written per second.
- ◆ **Avg Write Response Time** — The average time it takes the Symmetrix system to perform write operations.
- ◆ **Total Capacity** — The total available capacity in GBs for disks in the external disk group.
- ◆ **Used Capacity** — The total used capacity in GBs for disks in the external disk group.

- ◆ **Members** — The number of members in the external disk group.

Charts:

See [External disk group metrics](#) on page 498.

Front-end director Diagnostic view

Diagnostic view provides the following information about front-end directors:

Table view:

- ◆ **ID** — The ID of each front end director.
- ◆ **% Busy** — The percent of time the directory is busy.
- ◆ **Host IOs/sec** — The number of host commands for data transfer per second.
- ◆ **Host MBs/sec** — The size (MBs) of the host throughput per second.
- ◆ **Reqs/sec** — The number of data transfers between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
- ◆ **System WP Events/sec** — The number of write misses due to the system write pending limit having been reached.
- ◆ **Device WP Events/sec** — The number of write misses due to the volume write pending limit having been reached.

Charts:

See [Front-end director metrics](#) on page 517.

Back-end director Diagnostic view

Diagnostic view provides the following information about back-end directors (DA and DX):

Table view:

- ◆ **ID** — The ID of each back end director.
- ◆ **% Busy** — The percent of time that a director is busy.
- ◆ **IOs/sec** — The number of IOs (per second) commands to the disk.
- ◆ **Reqs/sec** — The number of data transfers between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. For writes the request counter is incremented at the time that the write pending flag is removed from the cache slot. In the event that multiple DAs are involved in the IO operation (such as RAID-1), the request count may not reconcile with the IO count and IO size.
- ◆ **MBs Read/sec** — The average size of host MBs read (per second) by the director.
- ◆ **MBs Written/sec** — The average size of host MBs written (per second) by the director.

Charts:

See [Back-end director metrics](#) on page 487.

Front-end ports Diagnostic view

Diagnostic view provides the following information about front-end ports:

Table view:

- ◆ **ID** — The port number.
- ◆ **Quota IOs/sec** — The number of IOs the FE port is completing for the specific SG.
- ◆ **Quota MBs/sec** — The number of MBs the FE port is processing for the specific SG.
- ◆ **IO Quota Limit** — The configured limit for IO quota (user-defined).
- ◆ **BW Quota Limit** — The configured limit for bandwidth quota (user-defined).
- ◆ **Quota Delayed IOs/sec** — The sum of the IOs that were delayed due to the SG's IO quota and bandwidth quota.
- ◆ **% Quota Delayed IOs** — The quota delayed IOs divided by the IO count.
- ◆ **Avg Quota Delayed Time** — The time that IOs are delayed (due to SG quotas) divided by the count of the IOs delayed (due to SG quotas).

Charts:

See [Front-end port metrics](#) on page 524.

RDF directors Diagnostic view

Diagnostic view provides the following information about RDF directors:

Table view:

- ◆ **ID** — The ID of each RDF director.
- ◆ **% Busy** — The percent of time that a director is busy.
- ◆ **IOs/sec** — The number of IOs (per second) handled by the director.
- ◆ **MBs Sent/sec** — The average size of MBs sent (per second) by the director.
- ◆ **MBs Received/sec** — The average size of MBs received (per second) by the director.
- ◆ **Copy IOs/sec** — The number of copy IOs (per second) by the director.
- ◆ **Copy MBs/sec** — The size of copy IOs (per second) by the director.
- ◆ **% Compressed MBs Sent** — The percent of total MBs sent by the director that were compressed MBs.

Charts:

See [RDF director metrics](#) on page 526.

RDF/S group Diagnostic view

Diagnostic view provides the following information about RDF/S groups:

Table view:

- ◆ **ID** — The RDF/S group number.
- ◆ **Response Time (ms)** — The average time it took the Symmetrix system to serve one IO for the RDF/S group.
- ◆ **Total HA Writes per second** — The total host writes per second for the RDF/S group.
- ◆ **% Hit** — The percent of read and write operations for this group that were immediately satisfied from cache.
- ◆ **% Write** — The percent of IO operations that were writes for the RDF/S group.

Charts:

See [Group metrics](#) on page 502.

RDF/A group Diagnostic view

Diagnostic view provides the following information about RDF/A groups:

Table view:

- ◆ **ID** — Shows the RDF/A group number.
- ◆ **Duration of Last Cycle** — The number of seconds for the last cycle in the session.
- ◆ **Host IOs/sec** — The total IOs per second for the RDF/A group.
- ◆ **RDF R1 to R2 MBs/sec** — The throughput in Mbytes per second from the R1 to the R2 for the RDF/A group.
- ◆ **RDF/A WP Count** — The number of cache slots in use by the RDF/A group.
- ◆ **Local WP Count** — The number of write pending slots waiting to be de-staged to disk. The R1 should be less than or equal to the system write pending count.
- ◆ **DSE Used Tracks** — The number of tracks used by DSE for the RDF/A group.

Charts:

See [RDF/A group metrics](#) on page 529.

Cache partition Diagnostic view

Diagnostic view provides the following information about cache partitions:

Table view:

- ◆ **ID** — The cache partition identifier.
- ◆ **% WP Utilization** — The percent of tracks in cache that are write pending. Until the data is destaged to disk, those tracks are not used to serve reads and improve the hit ratio.
- ◆ **% Hit** — The Read/Hit and Write/Hit ratio percentage for the cache partition.
- ◆ **% Cache Used** — The percent of the cache partition that is used.

Charts:

See [Cache partition metrics](#) on page 493.

Disk group tier Diagnostic view

Diagnostic view provides the following information about Symmetrix disk group tiers:

Table view:

- ◆ **ID** — The name assigned to this storage tier.
- ◆ **% Busy** — A high-level overview of the utilization (time busy) for the disks that comprise this storage tier during the selected time range.
- ◆ **Reads/sec** — The average number of host Reads performed per second.
- ◆ **MBs Read/sec** — The throughput in MBs read per second.
- ◆ **Avg Read Response Time** — The average time it takes the Symmetrix system to perform the Reads, in milliseconds.
- ◆ **Writes/sec** — The average number of host Writes performed per second.
- ◆ **MBs/Written/sec** — The throughput in MBs written per second.
- ◆ **Avg Write Response Time** — The average time it takes the Symmetrix system to perform the Writes, in milliseconds.
- ◆ **Total Capacity (GB)** — The total available capacity in GBs for disks that comprise this storage tier.
- ◆ **% Used Capacity** — The total capacity used in GBs for disks that comprise this storage tier.
- ◆ **Number of Disks**— The number of members in the tier.
- ◆ **Technology Protection** — The protection type of the tier.

Charts:

See [Disk group storage tier metrics](#) on page 514.

Virtual pool tier Diagnostic view

Diagnostic view provides the following information about Symmetrix virtual pool tiers:

Table view:

- ◆ **ID** — The name assigned to this storage tier.
- ◆ **BE Req/sec** — The number of read/write requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Transferred/sec** — The number of MBs per sec + the number of MBs written per second.
- ◆ **BE Avg Response Time (ms)** — The average time it takes to perform IO requests from the disk directors to cache.
- ◆ **Allocated Pool capacity (GB)** — The allocated pool capacity.

Charts:

See [Virtual pool tier metrics](#) on page 565.

Virtual pool Diagnostic view

Diagnostic view provides the following information about Symmetrix virtual pool tiers:

Table view:

- ◆ **ID** — The name assigned to this storage tier.
- ◆ **% Busy** — A high-level overview of the utilization (time busy) for the disks that comprise this storage tier during the selected time range.
- ◆ **Reads/sec** — The average number of Reads performed per second.
- ◆ **MBs Read/sec** — The average MBs Read per second.
- ◆ **Avg Read Time (ms)** — The average time to perform a Read in milliseconds.
- ◆ **Writes/sec** — The average number of Writes performed per second.
- ◆ **MBs Written/sec** — The average MBs Written per second.
- ◆ **Avg Write Time (ms)** — The average time to perform a Write in milliseconds.
- ◆ **Total Capacity (GB)** — The total available capacity in GBs for pools that comprise this thin tier.
- ◆ **Used Capacity (%)** — The total capacity used in GBs for pools that comprise this thin tier.

Charts:

See [Thin pool metrics](#) on page 553.

Volume Diagnostic view

Diagnostic view provides the information about volumes:

Table view:

- ◆ **ID** — The Symmetrix volume number.
- ◆ **Host IOs per second** — The number of host read IO and write IO operations performed each second by the Symmetrix volume.
- ◆ **Avg Read Time (ms)** — The average time it takes the volume to perform the Reads in milliseconds.
- ◆ **Avg Write Time (ms)** — The average time it takes the volume to perform the Writes in milliseconds.
- ◆ **% Hit** — The percentage of IO operations, performed by the Symmetrix volume, that were immediately satisfied by cache.
- ◆ **% Write** — The percent of IO operations that were writes.
- ◆ **% Read Miss** — The percent of read miss operations performed each second by the volume. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.
- ◆ **RAID Type** — The RAID protection type of the volume.

- ◆ **Capacity (GB)** — The capacity of the volume in GBs.
- ◆ **Number of Members** — The number of volumes.

Charts:

See [Volume metrics](#) on page 557.

Storage group Diagnostic view

Diagnostic view provides the following information about storage groups:

Table view:

- ◆ **ID** — The name assigned to this storage group.
- ◆ **Host IOs/sec** — The number of host read IO and write IO operations performed each second by the group.
- ◆ **Host MBs/sec** — The number MBs performed each second by the group.
- ◆ **Read Response Time (ms)** — The time it takes the group to perform the Reads in milliseconds.
- ◆ **Write Response Time (ms)** — The time it takes the group to perform the Writes in milliseconds.
- ◆ **% Hit** — The percentage of IO operations, performed by the group, that were immediately satisfied by cache.
- ◆ **% Writes** — The percent of IO operations that were writes.
- ◆ **% Read Miss** — The percent of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.
- ◆ **Capacity (GB)** — The capacity of the storage group in GBs.
- ◆ **Members** — The number of volumes that comprise this storage group.

Charts:

See [Storage group explore metrics](#) on page 533.

Snap pools Diagnostic view

Diagnostic view provides the following information about snap pools:

Table view:

- ◆ **ID** — The identifier of the snap pool.
- ◆ **BE Reqs/sec** — The number of read/write requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Reads/sec** — The throughput (MBs) that were reads.
- ◆ **BE MBs Written/sec** — The throughput (MBs) that were writes.
- ◆ **BE Disk Read Response Time (ms)** — The calculated back-end disk read response time in milliseconds for the snap pool.

- ◆ **Allocated Pool Capacity (GB)** — The allocated capacity for the snap pool in GBs.
- ◆ **Used Pool Capacity (GB)** — The used pool capacity in GBs.
- ◆ **Members** — The number of volumes in the snap pool.

Charts:

See [Snap and DSE pool metrics](#) on page 533.

SAVE volume Diagnostic view

Diagnostic view provides the following information about SAVE volumes:

Table view:

- ◆ **ID** — The Symmetrix volume number of the SAVE volume.
- ◆ **BE Requests/sec** — The total back-end director requests for this SAVE volume.
- ◆ **BE MBs Read/sec** — The number of MBs read by the disk directors from the disk each second.
- ◆ **BE MBs Written/sec** — The number of MBs written to the disk from the disk director each second.
- ◆ **BE Disk Read Response Time (ms)** — The calculated average of the read operations on the back-end disks.

Charts:

See [Volume metrics](#) on page 557.

Hyper Diagnostic view

Diagnostic view provides the following information about the related hypers:

Table view:

- ◆ **Hyper number** — The hyper number.
- ◆ **Symm Dev** — The Symmetrix volume number for the hyper.
- ◆ **Host IOs/sec** — The total number of host read IO and write IO operations performed each second by the Symmetrix volume.
- ◆ **BE % Read** — The percent of IO operations that were reads on the back-end.
- ◆ **BE % Write** — The percent of IO operations that were writes on the back-end.
- ◆ **BE Disk Response Time (ms)** — The calculated response time for IO operations on the back end.
- ◆ **Group** — The disk group number.
- ◆ **Type** — The type of hyper.

Charts:

See [Volume metrics](#) on page 557.

Thin pool Diagnostic view

Diagnostic view provides the following information about virtual pools:

Table view:

- ◆ **ID** — The identifier of the thin pool.
- ◆ **BE Reqs/sec** — The number of data transfers between cache and the pool.
- ◆ **BE MBs Transferred/sec** — The number of MBs transferred each second between cache and the pool.
- ◆ **BE Avg Response Time (ms)** — The average time it takes to transfer IO from pool to cache.
- ◆ **Allocated Pool Capacity (GB)** — The allocated capacity for the thin pool in GBs.

Charts:

See [Thin pool metrics](#) on page 553.

Thin volume Diagnostic view

Diagnostic view provides the following information about thin volumes:

Table view:

- ◆ **ID** — The Symmetrix volume number.
- ◆ **Host IOs/sec** — The number of host read IO and write IO operations performed each second by the Symmetrix volume.
- ◆ **BE Reqs/sec** — The number of IOs each second performed by the disk directors to cache.
- ◆ **Avg Read Time (ms)** — The average time it takes the volume to perform the Reads in milliseconds.
- ◆ **Avg Write Time (ms)** — The average time it takes the volume to perform the Writes in milliseconds.
- ◆ **% Hit** — The percentage of IO operations, performed by the Symmetrix volume, that were immediately satisfied by cache.
- ◆ **% Write** — The percent of IO operations that were writes.
- ◆ **% Read Miss** — The percent of read miss operations performed each second by the volume. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.
- ◆ **RAID Type** — The RAID protection type of the volume.
- ◆ **Capacity (GB)** — The capacity of the volume in GBs.
- ◆ **Members** — The number of volumes.

Charts:

See [Volume metrics](#) on page 557.

DATA volume Diagnostic view

Diagnostic view provides the following information about DATA volumes:

Table view:

- ◆ **ID** — The Symmetrix volume number.
- ◆ **BE Reqs/sec** — The number of IOs each second performed by the disk directors to cache.
- ◆ **BE MB Read/sec** — The size of the reads (MBs) each second performed by the disk directors to cache.
- ◆ **BE MB Written/sec** — The size of the writes (MBs) each second performed by the disk directors to cache.
- ◆ **BE Disk Read Response Time (ms)** — The calculated response time for read operations on the back end.
- ◆ **Allocated Capacity (GB)** — The allocated capacity for the DATA volume.
- ◆ **Used Capacity (GB)** — The total capacity used for the DATA volume.
- ◆ **Members** — The number of volumes members.

Charts:

See [Volume metrics](#) on page 557.

DSE pools Diagnostic view

Diagnostic view provides the following information about DSE pools:

Table view:

- ◆ **ID** — The identifier of the DSE pool.
- ◆ **BE Reqs/sec** — The number of read/write requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Read/sec** — The number of MBs that were read requests to the disk directors from the disk.
- ◆ **BE MBs Written/sec** — The number of MBs that were written to the disk directors from the disk.
- ◆ **BE Disk Read Response Time (ms)** — The calculated response time for read operations on the back end.
- ◆ **Allocated Pool Capacity (GB)** — The allocated capacity for the DSE pool in GBs.
- ◆ **Used Pool Capacity (GB)** — The used pool capacity in GBs.
- ◆ **Members** — The number of volumes in the DSE pool.

Charts:

See [Snap and DSE pool metrics](#) on page 533.

Disk Buckets diagnostic view

The Performance analyzer creates 10 buckets for each disk group. The buckets are based on disk utilization (% Busy) where each bucket summarizes (or averages) key metrics of disks that their utilization in the last sample put them in the bucket.

Table view:

- ◆ **ID** — The disk ID.
- ◆ **Avg Number Disks** — The number of disks in this bucket.
- ◆ **% Busy** — The percent of busy disks in this bucket.
- ◆ **Avg IOs/sec**— The average number of disk IOs of disks in this bucket.
- ◆ **Avg MBs/sec**— The average number of MBs read and written to disks in this bucket.
- ◆ **Avg Queue Depth** — The average queue depth of disks in this bucket.
- ◆ **Avg. Response Time** — The weighted average response time (reads and writes) for disks in this bucket.
- ◆ **Total Capacity (GB)** — The total capacity of disks in this bucket.
- ◆ **Used Capacity (%)** — The used capacity of disks in this bucket.

Charts:

See [Disk Buckets Explore metrics](#) on page 508.

Disk groups Diagnostic view

Diagnostic view provides the following information about disk groups:

Table view:

- ◆ **ID** — The disk group ID.
- ◆ **Name** — The disk group name.
- ◆ **% Busy** — The percent of time the disk group is busy.
- ◆ **Reads/sec** — The number of read operations performed by the disk group per second.
- ◆ **MBs Read/sec** — The number of read MBs performed by the disk group per second.
- ◆ **Avg Read Response Time (ms)** — The average time it took the disk group to serve one read command.
- ◆ **MBs Written/sec** — The number of write MBs performed by the disk group per second.
- ◆ **Avg Write Response Time (ms)** — The average time it took the disk group to serve one write command.
- ◆ **Total Capacity (GB)** — The total capacity of the disk group.
- ◆ **Used Capacity (%)** — The used capacity of the disk group.
- ◆ **Members** — The number of members in the disk group.

- ◆ **Model** — The vendor model number of the disk.

Charts:

See [Disk group metrics](#) on page 512.

Disk Diagnostic view

Table view:

- ◆ **ID** — The disk ID.
- ◆ **Name** — The disk name.
- ◆ **% Busy** — The percent of time that the disk is busy serving IOs.
- ◆ **IOs/sec** — The number of host IO commands per second for data transfer.
- ◆ **MBs/sec** — The throughput number (MBs) per second from the host.
- ◆ **Average Queue Depth** — The calculated value: (Accumulated queue depth / Total SCSI commands per second)
- ◆ **Reads/sec** — The number of read IOs performed by the disk per second.
- ◆ **Avg. Read Response Time (ms)** — The average time it took the disk to serve one read command.
- ◆ **Writes/sec** — The number of write IOs performed by the disk per second.
- ◆ **Avg. Write Response Time (ms)** — The average time it took the disk to serve one write command.
- ◆ **Total Capacity (GB)** — The total capacity of the disk.
- ◆ **Used Capacity (%)** — The used capacity of the disk.
- ◆ **Group** — The disk group number.
- ◆ **Model** — The vendor model number of the disk.

Charts:

See [Disk metrics](#) on page 510.

Events Diagnostic view

Diagnostic view provides the following information about events:

- ◆ **Sequence Number** — The number of the event.
- ◆ **Timestamp** — The date and time the event was generated.
- ◆ **Function Class** — The grouping of the event according to Unisphere categories.
- ◆ **Action Code** — The type of action that was taken.
- ◆ **Activity ID** — An ID assigned to the activity performed.
- ◆ **Message** — A text description of the action.

Host IO Limits by SG Diagnostic view

Diagnostic view provides the following information about Host IO Limits set for the storage group.

Table view:

- ◆ **ID** — The name of the storage group.
- ◆ **Host IO Limit IOs/sec** — Number of IOs the FE port is completing for the specific SG.
- ◆ **Host IO Limit MBs/sec** — Number of MBs the FE port is processing for the specific SG.
- ◆ **Host IO Limit % Delayed IOs** — The Host IO delayed IOs divided by the IO count.
- ◆ **Host IO Limit Avg Delayed Time (ms)** — The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).

Charts:

See [Host IO Limits by SG metrics](#) on page 568.

Host IO Limits by FE Diagnostic view

Diagnostic view provides the following information about Host IO Limits sorted by the front-end director.

Table view:

- ◆ **ID** — The name of the director.
- ◆ **Host IO Limit IOs/sec** — Number of IOs the FE port is completing for the specific SG.
- ◆ **Host IO Limit MBs/sec** — Number of MBs the FE port is processing for the specific SG.
- ◆ **Host IO Limit % Delayed IOs** — The Host IO delayed IOs divided by the IO count.
- ◆ **Host IO Limit Avg Delayed Time (ms)** — The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).

Charts:

See [Host IO Limits by FE metrics](#) on page 569.

Analyze Historical

External Disk Historical view

Table view:

- ◆ **ID** — The numeric identifier of the external disk.
- ◆ **Name** — The external disk name.
- ◆ **IOs/sec** — The number of host IO commands per second for data transfer.
- ◆ **MBs/sec** — The throughput number (MBs) per second from the host.
- ◆ **Reads/sec** — The number of read IOs performed by the disk per second.

- ◆ **Avg. Read Response Time (ms)** — The average time it took the disk to serve one read command.
- ◆ **Writes/sec** — The number of write IOs performed by the disk per second.
- ◆ **Avg. Write Response Time (ms)** — The average time it took the disk to serve one write command.
- ◆ **Total Capacity (GB)** — The total capacity of the external disk.
- ◆ **Used Capacity (%)** — The used capacity of the external disk.
- ◆ **Group** — The disk group number.
- ◆ **Members** — The number of members for the external disk.

Charts:

See [External disk metrics](#) on page 499.

External Disk Group Historical view

Diagnostic view provides the following information about Symmetrix external disk groups:

Table view:

- ◆ **ID** — The identifier assigned to this external disk group.
- ◆ **Name** — The name of the external disk group.
- ◆ **Reads/sec** — The average number of host reads performed per second.
- ◆ **MBs Read/sec** — The throughput per second for the disks that comprise this disk group.
- ◆ **Avg Read Response Time** — The average time it takes the Symmetrix system to perform read operations.
- ◆ **Writes/sec** — The average number of host writes performed per second.
- ◆ **MBs Written/sec** — The number of MBs written per second.
- ◆ **Avg Write Response Time** — The average time it takes the Symmetrix system to perform write operations.
- ◆ **Total Capacity** — The total available capacity in GBs for disks in the external disk group.
- ◆ **Used Capacity** — The total used capacity in GBs for disks in the external disk group.
- ◆ **Members** — The number of members in the external disk group.

Charts:

See [External disk group metrics](#) on page 498.

FAST policy Historical view

Historical view provides the following information about FAST policies:

Table view:

- ◆ **ID** — The policy ID.

- ◆ **Allocated SG Capacity (GB)** — The number of GBs of the storage group that are allocated to the FAST policy.
- ◆ **Total SG Capacity (GB)** — The total capacity of the storage group.
- ◆ **Allocated Pool Capacity (GB)** — The number of GBs of the virtual pool that are allocated to the FAST policy.
- ◆ **Total Pool Capacity (GB)** — The total capacity of the virtual pool.
- ◆ **BE Reqs/sec** — The number of data transfers between cache and the director.
- ◆ **BE MBs Transferred/sec** — The number of MBs transferred each second between cache and the director.

Charts:

See [FAST policy metrics](#) on page 516.

FE Port SG Historical view

Historical view provides the following information about front-end quotas sorted by the front-end port and storage group.

Table view:

- ◆ **ID** — The name of the storage group or port.
- ◆ **Host IO Limit IOs/sec**— The user-defined Host IO limit data transferred per second.
- ◆ **Host IO Limit MBs/sec** — The host throughput within the limit set for the specific SG.
- ◆ **Host IO Limit % Delayed IOs** — The percent of delayed IOs divided by the delayed IO count.
- ◆ **Host IO Limit Avg Delayed Time** —The average time the data transferred was delayed by Host IO Limits.

Charts:

See [FE Port SG metrics](#) on page 524.

Symmetrix system Historical view

Historical view provides the following information about Symmetrix systems:

Table view:

- ◆ **Symmetrix ID** — The Symmetrix identifier.
- ◆ **Host IOs/sec** — Number of host IO operations performed each second by all Symmetrix volumes, including writes and random and sequential reads.
- ◆ **Host MBs/sec** — The number of host MBs written and read by all of the Symmetrix volumes each second.
- ◆ **FE Utilization** — The percent of time the front-end director is busy.
- ◆ **BE Utilization** — The percent of time the back-end director is busy.
- ◆ **RDF Utilization** — The percent of time the RDF director is busy.

- ◆ **% Cache WP** — The percent of system cache that is write pending.
- ◆ **Disk Utilization** — The percent of time the disks are busy.
- ◆ **Avg Fall Thru Time** — Average time it takes a cache slot to be freed up. It is the average time from the first use of the contents to its reuse by another address.

Charts:

See [Symmetrix system Historical metrics](#) on page 542.

Storage groups Historical view

Historical view provides the following information about storage groups:

Table view:

- ◆ **ID** — The storage group identifier.
- ◆ **Host IOs/sec** — The number of IOs per second transferred from the host.
- ◆ **Host MBs/sec** — Number MBs performed each second by the storage group.
- ◆ **Read Response Time (ms)** — The average time it takes the Symmetrix system to read IO for this group.
- ◆ **Write Response Time (ms)** — The average time it takes the Symmetrix system to write IO for this group.
- ◆ **% Hit** — The percent of IOs performed by the group that were immediately satisfied by cache.
- ◆ **% Writes** — The percent of IO operations that were writes.
- ◆ **% Read Miss** — Percent of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
- ◆ **Capacity (GB)** — The capacity of the storage group in GBs.
- ◆ **Members** — The number of volumes that comprise this group.

Charts:

See [Storage group explore metrics](#) on page 533.

RDF/A group Historical view

Historical view provides the following information about RDF/A groups:

Table view:

- ◆ **ID** — shows RDF/A group number.
- ◆ **Duration of Last Cycle** — The number of seconds for the last cycle in the session.
- ◆ **Host IOs/sec** — The total IOs per second for the RDF/A group.
- ◆ **RDF R1 to R2 MBs/sec** — The throughput in Mbytes per second from the R1 to the R2 for the RDF/A group.
- ◆ **RDF/A WP Count** — The number of cache slots in use by the RDF/A group.

- ◆ **Local WP Count** — The number of write pending slots waiting to be de-staged to disk. The R1 should be less than or equal to the system write pending count.
- ◆ **DSE Used Tracks** — The number of tracks used by DSE for the RDF/A group.

Charts:

See [Group metrics](#) on page 502.

RDF/S group Historical view

Historical view provides the following information about RDF/S groups:

Table view:

- ◆ **ID** — The RDF/S group number.
- ◆ **Response Time (ms)** — The average time it took the Symmetrix system to serve one IO for the RDF/S group.
- ◆ **Total HA Writes per second** — The total host writes per second for the RDF/S group.
- ◆ **% Hit** — The percent of read and write operations for this group that were immediately satisfied from cache.
- ◆ **% Write** — The percent of IO operations that were writes for the RDF/S group.

Charts:

See [Group metrics](#) on page 502.

FAST policy Historical view

Historical view provides the following information about FAST policies:

Table view:

- ◆ **ID** — The policy ID.
- ◆ **Allocated SG Capacity (GB)** — The number of GBs of the storage group that are allocated to the FAST policy.
- ◆ **Total SG Capacity (GB)** — The total capacity of the storage group.
- ◆ **Allocated Pool Capacity (GB)** — The number of GBs of the virtual pool that are allocated to the FAST policy.
- ◆ **Total Pool Capacity (GB)** — The total capacity of the virtual pool.
- ◆ **BE Reqs/sec** — The number of data transfers between cache and the director.
- ◆ **BE MBs Transferred/sec** — The number of MBs transferred each second between cache and the director.

Charts:

See [FAST policy metrics](#) on page 516.

Tier by Storage Group Historical view

Historical view provides the following information about tiers by storage group:

Table view:

- ◆ **ID** — The storage group name.
- ◆ **BE Read Reqs/sec** — The number of read requests each second performed by the disk directors to cache.
- ◆ **BE Write Reqs/sec** — The number of write requests each second performed by the disk directors to cache.
- ◆ **BE Reqs/sec** — The number of requests each second performed by the disk directors to cache.
- ◆ **BE MBs Read/sec** — The number of host read IOs performed by the disk group per second.
- ◆ **BE MBs Written/sec** — The number of host write IOs performed by the disk group per second.
- ◆ **BE MBs Transferred/sec** — The number of MBs performed by the disk group per second.
- ◆ **Allocated Capacity** — The number of GBs of the storage group that are allocated to the FAST policy.
- ◆ **IO Density** — The number BE requests per GB of disk. (BE Reads + destaged writes) / capacity

Charts:

See [Tier by Storage Group metrics](#) on page 551.

Front-end director Historical view

Historical view provides the following information about front-end directors:

Tableview:

- ◆ **ID** — The ID of the director.
- ◆ **% Busy** — The percent of time the directory is busy.
- ◆ **Host IOs/sec** — The IOs per second for each front end director.
- ◆ **Host MBs/sec** — The MBs per second for each front end director.
- ◆ **Reqs/sec** — The number of requests per second for each front end director.
- ◆ **System WP Events/sec** — The number of write pending events per second for the system.
- ◆ **Device WP Events/sec** — The number of write pending events per second for the volume.

Charts:

See [Front-end director metrics](#) on page 517.

Back-end director Historical view

Historical view provides the following information about back-end directors:

Table view:

- ◆ **ID** — The ID of each back-end director.
- ◆ **% Busy** — The percent of time that a director is busy.

- ◆ **IOs/sec** — The number of IOs (per second) commands to the disk.
- ◆ **Reqs/sec** — The number of data transfers between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. For writes the request counter is incremented at the time that the write pending flag is removed from the cache slot. In the event that multiple DAs are involved in the IO operation (such as RAID-1), the request count may not reconcile with the IO count and IO size.
- ◆ **MBs Read/sec** — The average size of MBs read (per second) by the director.
- ◆ **MBs Written/sec** — The average size of MBs written (per second) by the director.

Charts:

See [Back-end director metrics](#) on page 487.

RDF director Historical view

Historical view provides the following information about RDF directors:

Table view:

- ◆ **ID** — The ID of each back-end director.
- ◆ **% Busy** — The percent of time that a director is busy.
- ◆ **IOs/sec** — The number of IOs (per second) commands to the disk.
- ◆ **MBs Sent/sec** — The average size of MBs sent (per second) by the director.
- ◆ **MBs Received/sec** — The average size of MBs received (per second) by the director.

Charts:

See [RDF director metrics](#) on page 526.

Cache partition Historical view

Historical view provides the following information about cache partitions:

Table view:

- ◆ **ID** — The cache partition identifier.
- ◆ **% WP Utilization** — The percent of tracks in cache that are write pending. Until the data is destaged to disk, those tracks are not used to serve reads and improve the hit ratio.
- ◆ **% Hit** — The Read/Hit and Write/Hit ratio percentage for the cache partition.
- ◆ **% Cache Used** — The percent of the cache partition that is used.

Charts:

See [Cache partition metrics](#) on page 493.

Disk group Historical view

Historical view provides the following information about disk groups:

Table view:

- ◆ **ID** — The disk ID.

- ◆ **% Busy** — The percent of time that the disk is busy serving IOs.
- ◆ **Reads/sec** — The number of read IOs performed by the disk per second.
- ◆ **MBs Read/sec** — The number of MBs read by the disk per second.
- ◆ **Avg. Read Response Time (ms)** — The average time it took the disk to serve one read command.
- ◆ **Writes/sec** — The number of write IOs performed by the disk per second.
- ◆ **Avg. Write Response Time (ms)** — The average time it took the disk to serve one write command.
- ◆ **Total Capacity (GB)** — The total capacity of the disk.
- ◆ **% Used Capacity** — The used capacity of the disk.
- ◆ **Members** — The number of disk group members.
- ◆ **Model** — The vendor model number of the disk.

Charts:

See [Disk group metrics](#) on page 512.

Virtual pool tier Historical view

Historical view provides the following information about virtual pool tiers:

Table view:

- ◆ **ID** — The name assigned to this storage tier.
- ◆ **BE Req/sec** — The number of read/write requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Transferred/sec** — The size of the throughput transferred by the disk directors to the cache.
- ◆ **BE Avg Response Time (ms)** — The average response time for the back end director.
- ◆ **Allocated Pool Capacity** — The total allocated capacity for the virtual pool tier.

Charts:

See [Virtual pool tier metrics](#) on page 565.

Snap pool Historical view

Historical view provides the following information about snap pools:

Table view:

- ◆ **ID** — The identifier of the snap pool.
- ◆ **BE Reqs/sec** — The number of requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Read/sec** — The number of MBs read by the disk director from the disk each second.

- ◆ **BE MBs Written/sec** — The number of MBs written to the disk from the disk director each second.
- ◆ **BE Disk Read Response Time** — The average time it takes Read requests from the disk directors to cache.
- ◆ **Allocated Pool Capacity (GB)** — The allocated capacity for the snap pool in GBs.
- ◆ **Used Pool Capacity (GB)** — The used pool capacity in GBs.
- ◆ **Members** — The number of volumes in the snap pool.

Charts:

See [Snap and DSE pool metrics](#) on page 533.

Thin pool Historical view

Historical view provides the following information about thin pools:

Table view:

- ◆ **ID** — The identifier of the thin pool.
- ◆ **BE Req/sec** — The number of read/write requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Transferred/sec** — The size of the throughput transferred by the disk directors to the cache.
- ◆ **BE Avg Response Time (ms)** — The average response time for the back end director.
- ◆ **Allocated Pool Capacity** — The total allocated capacity for the thin pool.

Charts:

See [Thin pool metrics](#) on page 553.

DSE pool Historical view

Historical view provides the following information about DSE pools:

Table view:

- ◆ **ID** — The identifier of the DSE pool.
- ◆ **BE Reqs/sec** — The number of read/write requests each second performed by the disk directors to the cache.
- ◆ **BE MBs Read/sec** — The number of MBs that were read requests to the disk directors from the disk.
- ◆ **BE Disk Read Response Time (ms)** — The calculated response time for read operations on the back end.
- ◆ **Allocated Pool Capacity (MB)** — The allocated capacity for the DSE pool in MBs.
- ◆ **Used Pool Capacity (MB)** — The used pool capacity in MBs.
- ◆ **Members** — The number of volumes in the DSE pool.

Charts:

See [Snap and DSE pool metrics](#) on page 533.

Disk Buckets historical view

The Performance analyzer creates 10 buckets for each disk group. The buckets are based on disk utilization (% Busy) where each bucket summarizes (or averages) key metrics of disks that their utilization in the last sample put them in the bucket.

Table view:

- ◆ **ID** — The disk ID.
- ◆ **Avg Number Disks** — The number of disks in this bucket.
- ◆ **% Busy** — The percent of busy disks in this bucket.
- ◆ **Avg IOs/sec**— The average number of disk IOs of disks in this bucket.
- ◆ **Avg MBs/sec**— The average number of MBs read and written to disks in this bucket.
- ◆ **Avg Queue Depth** — The average queue depth of disks in this bucket.
- ◆ **Avg. Response Time** — The weighted average response time (reads and writes) for disks in this bucket.
- ◆ **Total Capacity (GB)** — The total capacity of disks in this bucket.
- ◆ **Used Capacity (%)** — The used capacity of disks in this bucket.

Charts:

See [Disk Buckets Explore metrics](#) on page 508.

Disk group tier Historical view

Historical view provides the following information about disk group tiers:

Table view:

- ◆ **ID** — The name assigned to this storage tier.
- ◆ **% Busy** — A high-level overview of the utilization (time busy) for the disks that comprise this storage tier during the selected time range.
- ◆ **Reads/sec** — The average number of Reads performed per second.
- ◆ **MBs Read/sec**— The average size of the Reads performed per second.
- ◆ **Avg Read Response Time (ms)** — The average time it takes the Symmetrix system to perform the Reads, in milliseconds.
- ◆ **Writes/sec** — The average number of Writes performed per second.
- ◆ **Avg Write Response Time (ms)** — The average time it takes the Symmetrix system to perform the Writes, in milliseconds.

- ◆ **Total Capacity (GB)** — The total available capacity in GBs for disks that comprise this storage tier.
- ◆ **% Used Capacity** — The total capacity used in GBs for disks that comprise this storage tier.
- ◆ **Number of Disks** — The number of disks in the tier.
- ◆ **Technology Protection** — The technology type (FC, EFD, SATA) and the RAID protection.

Charts:

See [Disk group storage tier metrics](#) on page 514.

Host IO Limits by SG Historical view

Historical view provides the following information about Host IO Limits:

Table view:

- ◆ **ID** — The name of the storage group.
- ◆ **Host IOs/sec** — The number of host commands for data transfer per second.
- ◆ **Read Response Time (ms)** — The average time it takes the director/port to service one Read IO for the SG.
- ◆ **Write Response Time (ms)** — The average time it takes the director/port to service one Write IO for the SG.
- ◆ **IO Limit % Delayed IOs** — The Host IO delayed IOs divided by the IO count.
- ◆ **Host IO Limit Avg Delayed Time (ms)** — The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
- ◆ **Host IO Limit Delayed IOs/sec** — The sum of the IOs that were delayed due to the SG's IO limits and bandwidth limits.

Charts:

See [Host IO Limits by SG metrics](#) on page 568.

Host IO Limits by FE Historical view

Diagnostic view provides the following information about Host IO Limits sorted by the front-end director.

Table view:

- ◆ **ID** — The name of the director.
- ◆ **Host IO Limit IOs/sec** — Number of IOs the FE port is completing for the specific director.
- ◆ **Host IO Limit MBs/sec** — Number of MBs the FE port is processing for the specific director.
- ◆ **Host IO Limit % Delayed IOs** — The Host IO delayed IOs divided by the IO count.

- ◆ **Host IO Limit Avg Delayed Time (ms)** — The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).

Charts:

See [Host IO Limits by FE metrics](#) on page 569.

Metric Tables

Symmetrix systems metrics

This is a list of all KPIs that are available for all Symmetrix systems.

Table 10: Symmetrix system metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second by all Symmetrix volumes, including writes and random and sequential reads.
Host Reads/sec	The number of host read operations performed each second by all Symmetrix volumes.
Host Writes/sec	The number of host write operations performed each second by all Symmetrix volumes.
Host MBs/sec	The number of host MBs written and read by all of the Symmetrix volumes each second.
Host MBs Read/sec	The number of host MBs read by all of the Symmetrix volumes each second.
Host MBs Written/sec	The number of host MBs written by all of the Symmetrix volumes each second.
FE Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
Avg Read Response Time (ms)	The average of the total read response time for all front-end directors in milliseconds.
Avg Write Response Time (ms)	The average of the total write response time for all front-end directors in milliseconds.
% Reads	The percent of total read IO operations performed each second by all of the Symmetrix volumes.
% Writes	The percent of total write IO operations performed by all of the Symmetrix volumes.
% Hit	The percent of IO operations performed by all of the Symmetrix volumes, for

Table 10: Symmetrix system metrics (continued)

Metric	Description
	which the read data was in cache and the write operation could be sent directly to cache without having to wait for data to be destaged from cache to the disks.
FE Read Reqs/sec	A read data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
FE Write Reqs/sec	A write data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
BE IOs/sec	The total IO from all BE directors to the disks per second.
BE Reqs/sec	A data transfer of a read or write between the cache and the director.
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
BE Write Reqs/sec	A data transfer of a write between the cache and the director.
System WP Events/sec	The number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific Symmetrix volume was reached.
WP Count	The number of system cache slots that are write pending.
System Max WP limit	The percent of the target % at which writes are delayed. The range is from 40% to 80%.
% Cache WP	The percent of system cache that is write pending.
Avg Fall Thru Time	The average time it takes a cache slot in LRU0 to be freed up. It is the average time from the first use of the contents to its reuse by another address.
FE Hit Reqs/sec	The total requests from all front-end directors per second that were satisfied from cache.
FE Read Hit Reqs/sec	The total read requests from all front-end directors per second that were satisfied from cache.
FE Write Hit Reqs/sec	The total write requests from all front-end directors per second that were satisfied from cache.
Prefetched Tracks/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.

Table 10: Symmetrix system metrics (continued)

Metric	Description
Destaged Tracks/sec	The number of tracks per second saved into disks.
FE Read Miss Reqs/sec	The total read requests from all front-end directors per second that were misses. A miss occurs when the requested data is not in cache.
FE Write Miss Reqs/sec	The total write requests from all front-end directors per second that were misses. A miss occurs when the write had to wait while data was destaged from cache to the disks.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read for this volume.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write for this volume.
XtremSW Cache IOs/sec	The average number of IOs served by XtremSW Cache per second.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache	The number of XtremSW Cache Dedup writes per second.

Table 10: Symmetrix system metrics (continued)

Metric	Description
Dedup Writes/sec	
XtremSW Cache Reads	The number of read operations served by XtremSW Cache per second.
XtremSW Cache Writes	The number of write operations served by XtremSW Cache per second.
XtremSW Cache Read Hits	The number of read operations that were hits served by XtremSW Cache per second.
XtremSW Cache MBs Read/sec	Cumulative number of host MBs read by the XtremSW Cache per second.
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache Read Time	The average time it took the XtremSW Cache to serve one read for this volume.
XtremSW Cache Write Time	The average time it took the XtremSW Cache to serve one read for this volume.

Back-end director metrics

This is a list of all KPIs that are available for back-end directors.

Table 11: Back-end director metrics

Metric	Description
% Busy	The percent of time that a director is busy.
% Idle	The percent of time that a director is idle.
IOs/sec	The number of IOs performed each second by the director.
Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. For writes the request counter increments at the time that the write pending flag is removed from the cache slot. In the event that multiple DAs are involved in the IO operation (such as RAID-1), the request count may not reconcile with the IO count and IO size.
Read Reqs/sec	A data transfer of a read between the director and the cache.
Write Reqs/sec	A data transfer of a write between the cache and the director.
MBs/sec	The total IO (reads and writes) per second in MBs.
Port 0 MBs/sec	The total MBs moving through port 0 per second.
Port 0 IOs/sec	The number of IOs moving through port 0 per second.
Port 1 MBs/sec	The total MBs moving through port 1 per second.
Port 1 IOs/sec	The number of IOs moving through port 0 per second.
% Busy Port 0	The utilization of port 0.
% Busy Port 1	The utilization of port 1.
% Read Reqs	The percent of read requests out of the total requests.
% Write Reqs	The percent of write requests out of the total requests.
Syscall Count/sec	The total number of calls seen by this director.
Avg Time per Syscall	The average time spent processing all system calls.

Table 11: Back-end director metrics (continued)

Metric	Description
Syscall Remote Dir Count/sec	The number of calls sent from the local director to another director in the same system.
Syscall RDF Dir Count/sec	The number of calls sent using RDF to a remote system.
Prefetched Tracks/sec	The number of tracks pre-fetched from disk to cache upon detection of a sequential read stream.
MBs Read/sec	The number of reads per second in MBs.
MBs Written/sec	The number of writes per second in MBs.
Port 0 Avg Req Size (KB)	The average IO request moving through port 0 per second.
Port 1 Avg Req Size (KB)	The average IO request moving through port 1 per second.
Copy On First Access	These metrics are for internal Symmetrix operations.
VLUN Migration Read	
VLUN Migration Write	
PHCO Rebuild Copy	
Optimized Write	
Port 0 Speed Gb/sec	The average IO request moving through port 0 per second.
Port 1 Speed Gb/sec	The average IO request moving through port 1 per second.

Table 11: Back-end director metrics (continued)

Metric	Description
PHCO Rebuild Copy	
Compressed Read Reqs/sec	The number of compressed read requests per second.
Compressed Write Reqs/sec	The number of compressed write requests per second.
Compressed Read MBs/sec	The size of compressed read MBs per second.
Compressed Write MBs/sec	The size of compressed write MBs per second.
Compressed Reqs/sec	The total number of compressed requests (read and write) per second.
Compressed MBs/sec	The total size of compressed MBs (read and write) per second.
% Compressed Reqs	The total percent of all read and write requests.
% Compressed Read MBs	The percent of all compressed MBs that were read requests.
% Compressed Write MBs	The percent of all compressed MBs that were write requests.
% Compressed MBs	The total percent of all read and write compressed MBs.
% Busy Logical Core 0	The percent of time that core 0 is busy serving IOs.
% Busy Logical Core 1	The percent of time that core 1 is busy serving IOs.
% Idle	The percent of time that core 0 is idle.

Table 11: Back-end director metrics (continued)

Metric	Description
Logical Core 0	
% Idle Logical Core 1	The percent of time that core 1 is idle.
Port 0 Reads/sec	The count of reads passed through port 0 per second.
Port 1 Reads/sec	The count of reads passed through port 1 per second.
Port 0 Writes/sec	The count of writes passed through port 0 per second.
Port 1 Writes/sec	The count of writes passed through port 1 per second.
Port 0 MBs Read/sec	The number of MBs served by port 0 per second.
Port 1 MBs Read/sec	The number of MBs served by port 1 per second.
Port 0 MBs Written/sec	The number of MBs per second on port 0 that were writes.
Port 1 MBs Written/sec	The number of MBs per second on port 1 that were writes.
% Non-IO Busy Logical Core 0	The percent of time that core 0 is busy serving other non-IO tasks, such as FAST swaps, Optimizer moves, Snap copies, etc.
% Non-IO Busy Logical Core 1	The percent of time that core 1 is busy serving other non-IO tasks, such as FAST swaps, Optimizer moves, Snap copies, etc.
% Non-IO Busy	The percent of time the director was busy with non-IO tasks, such as FAST swaps, Optimizer moves, Snap copies, etc.

Back-end director Real Time metrics

This is a list of all KPIs that are available in Real Time view for back-end directors (DA and DX).

Table 12: Back-end director Real Time metrics

Metric	Description
% Busy	The percent of time that a director is busy.
IOs/sec	The number of IO commands to the disk per second.
Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. For writes the request counter increments at the time that the write pending flag is removed from the cache slot. In the event that multiple DAs are involved in the IO operation (such as RAID-1), the request count may not reconcile with the IO count and IO size.
Read Reqs/sec	A data transfer of a read between the director and the cache.
Write Reqs/sec	A data transfer of a write between the cache and the director.
Port 0 IOs/sec	The count of the IOs passed through port 0 per second.
Port 0 MBs/sec	The size of the IOs in MBs passed through port 0 per second.
Port 1 IOs/sec	The count of the IOs passed through port 1 per second.
Port 1 MBs/sec	The size of the IOs in MBs passed through port 1 per second.
Port 0 Reads/sec	The number of read operations that passed through port 0 per second.
Port 1 Reads/sec	The number of read operations that passed through port 1 per second.
Port 0 Writes/sec	The number of write operations that passed through port 0 per second.
Port 1 Writes/sec	The number of write operations that passed through port 1 per second.
Port 0 MBs Read/sec	The number of MBs per second on port 0 that were reads.
Port 1 MBs Read/sec	The number of MBs per second on port 1 that were reads.
Port 0 MBs Written/sec	The number of MBs per second on port 0 that were writes.

Table 12: Back-end director Real Time metrics (continued)

Metric	Description
Port 1 MBs Written/sec	The number of MBs per second on port 1 that were writes.

Cache partition metrics

This is a list of all KPIs that are available for cache partitions.

Table 13: Cache partition metrics

Metric	Description
WP Count	A sampled value at the time of data collection.
WP Limit	The percent of the target % at which writes are delayed The range is from 40% to 80%. If the target is 30% and the write pending limit is 80%, then the actual limit is 24% of the cache.
% Target Cache	The percent of total cache allocated to this partition; cannot be less than 10% or larger than 90% and the sum of all targets must be equal to 100%.
% Min Cache	The partition will not go below this allocation (from 0% to target %).
% Max Cache	The maximum slot allocation for a partition (from target % to 100%).
Cache Slots Used	A sampled value at the time of data collection.
Donation Time	The length of time a cache slot must age before being available for donation.
% WP Utilization	A calculated value $(WP\ Count / (total\ slots * \% \text{ Cache Used} / 100) * 100)$
% Hit	The percent of IOs that were immediately satisfied.
% Cache Used	The percent of cache that is used.
Cache Age GT 1 Min	The average fall-through-time w/ 64 second decay. Decay is a moving average, where the decay time is the "average time" average.
Cache Age GT 10 Min	The average fall-through-time w/ 8:32 second decay. Decay is a moving average, where the decay time is the "average time" average.
Cache Age GT 1 Hour	The average fall-through-time w/ 1:08:16 second decay. Decay is a moving average, where the decay time is the "average time" average.
Cache Hit Ages 1 to Cache Hit Ages 8	A partition read hit under 4 seconds. A slot in this age range satisfied a host read request (cache hit) and was promoted to age zero.
Cache Slot Ages 1 to Cache Slot Ages 8	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Destage Slot Age	The average time accumulated for destaged cache slots.
Age Non WP Slots	The average time accumulated for non WP cache slots.

Table 13: Cache partition metrics (continued)

Metric	Description
Total Replace Slots	The partition fall-through-time events (FT).
Writes to Young Non WP Slots	The number of writes to cache slots that were less than four seconds old.
Writes to All Non WP Slots	The number of writes to all non WP cache slots that were 4 seconds or older.
Acquired Slot Count	The number of cache partition slots acquired from other cache.
Donation Give Count	The amount of cache donated from other cache.
Donation Take Count	The amount of cache donated to other cache.
Destaged Slot Count	The number of cache slots destaged.
Avg Age Given Destage	Calculated value: (destage slot age) / (destaged slot count).
Ave Age of Write to Non WP Slot	Calculated value: (age non WP slots) / ((writes to all non WP slots) - (writes to young non WP slots)).
Remote WP Count	The number of writes pending for a remote system.
Local WP Count	The calculation of the total writes pending minus the remote writes pending.
Host IOs/sec	The number of host IOs per second.
Host MBs/sec	The number of host MBs per second.

Device pool metrics

This is a list of all KPIs that are available for all volume pools (Snap, DSE, Thin).

Table 14: Device pool metrics

New Name	Description
Host IOs/sec	The number of host operations performed each second by the pool.
Host Reads/sec	The number of host read operations performed each second by the pool.
Host Writes/sec	The number of host write operations performed each second by the pool.
Host Hits/sec	The number of host read/write operations performed each second by the pool that were immediately satisfied from cache.
Host Read Hits/sec	The number of host read operations performed each second by the pool that were immediately satisfied from cache.
Host Write Hits/sec	The number of host write operations performed each second by the pool that were immediately satisfied from cache.
Host Misses/sec	The number of host read/write operations performed each second by the pool that could not be satisfied from cache.
Host Read Misses/sec	The number of host read operations performed each second by the pool that were not satisfied from cache.
Host Write Misses/sec	The number of host write operations performed each second by the pool that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the pool.
Host MBs Read/sec	The cumulative number of host MBs read per second by the pool.
Host MBs written/sec	The cumulative number of host MBs written per second by the pool.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
Read Response-Time (ms)	The average time that it took the Symmetrix to serve one read IO for this pool.
Write Response-Time	The average time that it took the Symmetrix to serve one write IO for this pool.

Table 14: Device pool metrics (continued)

New Name	Description
(ms)	
Read Miss Response Time (ms)	The average time that it took the Symmetrix to serve one read miss IO for this pool.
Write Miss Response Time (ms)	The average time that it took the Symmetrix to serve one write miss IO for this pool.
RDF/S Write Response Time (ms)	The average time that it took the Symmetrix RDF/S to serve one write IO for this pool.
% Read	The percent of IO operations that were reads.
% Write	The percent of IO operations that were writes.
% Read Hit	The percent of read operations, performed by the pool, that were immediately satisfied by cache.
% Write Hit	The percent of write operations, performed by the pool, that were immediately satisfied by cache.
% Read Miss	The percent of read miss operations performed each second by the pool. A miss occurs when the requested read data is not found in cache.
% Write Miss	The percent of write miss operations performed each second by the pool. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the pool
Seq IOS/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the pool that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the pool that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the pool that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the pool that were misses.

Table 14: Device pool metrics (continued)

New Name	Description
Random IOs/sec	The number of IOs from a host not identified as part of a sequential stream.
Random Reads/sec	The number of read IO commands from a host not identified as part of a sequential stream.
Random Writes/sec	The number of write IO commands from a host not identified as part of a sequential stream.
Random Read Hits/sec	The number of random read IOs that were satisfied from the cache.
Random Write Hits/sec	The number of random write IOs that were immediately placed in cache because space was available.
Random Read Misses/sec	The number of random read IOs that were misses.
Random Write Misses/sec	The number of random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Seq IO	Calculated value: $((100 * (\text{total seq IOs per sec} / \text{total IOs per sec}))$
% Seq Read	Calculated value: $((100 * (\text{seq reads per sec} / \text{total IOs per sec}))$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $(100 * (\text{seq writes per sec} / \text{total IOs per sec}))$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO that were not identified as part of a sequential stream.
% Random Read Hit	Calculated value: $((100 * (\text{random read hits per sec} / \text{total IOs per sec}))$
% Random	Calculated value: $((100 * (\text{random read misses per sec} / \text{total IOs per sec}))$

Table 14: Device pool metrics (continued)

New Name	Description
Read Miss	
% Random Write Hit	Calculated value: $((100 * (\text{random write hits per sec} / \text{total IOs per sec}))$
% Random Write Miss	Calculated value: $((100 * (\text{random write misses per sec} / \text{total IOs per sec}))$
Max WP Threshold	The maximum number of write-pending slots available for the pool.
BE MBs transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Request Time (ms)	The average time it takes to make a request by the disk directors to cache.
BE Disk Response-Time (ms)	The average time it takes cache to satisfy IO from the disk directors.
BE Task Time (ms)	The time from the point when the HA puts the request on the queue and the DA picks it up; this can be considered queue time.

External disk group metrics

This is a list of all KPIs that are available for external disk groups.

Table 15: External disk group metrics

Metric	Description
Reads/sec	The number of reads per second for the disk group.
Writes/sec	The number of writes per second for the disk group.
MBs/sec	The total number of MBs per second for the disk group.
MBs Read/sec	The read throughput (MBs) of the disk group per second.

Table 15: External disk group metrics (continued)

Metric	Description
MBs Written/sec	The write throughput (MBs) of the disk group per second.
Avg Read Size (KB)	The average number of kilobytes for a single read command.
Avg Write Size (KB)	The average number of kilobytes for a single write command.
Avg Response Time	The average time it took the disk group to service IO.
Avg Read Response Time (ms)	The average time it took the disk group to serve one read command.
Avg Write Response Time (ms)	The average time it took the disk group to serve one write command.
Total Capacity (GB)	The total capacity of all the disks in the disk group.
MBs/sec	The total number of MBs per second for the disk group.
IOs/sec	The total number of read and write IOs per second.
Used Capacity (GB)	The total capacity allocated for all the disks in the disk group.
% Used Capacity	The percent of the disk group capacity that is allocated.
Members	The number of disks in the disk group.
Avg Queue Depth	Calculated value: Accumulated queue depth/total SCSI command per sec.
% Idle	The percent of time the disk group is idle.
% Busy	The percent of time the disk group is busy.

External disk metrics

This is a list of available metrics for external disks.

Table 16: External disk metrics

New Name	Description
Reads/sec	The number of host reads per second for the external disk.
Writes/sec	The number of host writes per second for the external disk.
MBs/sec	The size of the IO from the host to the external disk per second.
MBs Read/sec	The read throughput (MBs) of the external disk per second.
MBs Written/sec	The write throughput (MBs) of the external disk per second.
Avg Read Size (KB)	The average number of kilobytes for a single read command.
Avg Write Size (KB)	The average number of kilobytes for a single write command.

Table 16: External disk metrics (continued)

New Name	Description
Avg Response Time	The average response time for the reads and writes.
Avg Read Response Time (ms)	The average time it took the external disk to serve one read command.
Avg Write Response Time (ms)	The average time it took the external disk to serve one write command.
MBs/sec	The total number of MBs per second for the external disk.
IOs/sec	The total number of read and write IOs per second for the external disk.
Total Capacity (GB)	The total capacity of the external disk (GBs).
Used Capacity (GB)	The used capacity of the external disk (GBs).
% Used Capacity	The percent of the external disk that is used.

Front-end port metrics

This is a list of all KPIs that are available for front-end ports that are associated with a storage group with Host IO Limits set on the directors.

Table 17: Front-end ports metrics

Metric	Description
Host IOs/sec	The number of IOs the port is performing on behalf of the specific SG.
IO Limit IOs/sec	The sum of the IOs that were delayed due to the SG's IO limits.
BW Limit Delayed IOs/sec	The sum of the MBs that were delayed due to the SG's bandwidth limits.
Host MBs/sec	The number of MBs the port is processing on behalf of the specific SG.
Host Limit Delayed IOs/sec	The sum of the IOs that were delayed due to the SG's IO limits and bandwidth limits.
% Host Limit Delayed IOs	The host limit delayed IOs divided by the IO count.
% IO Host Limit Delayed IOs	The IO delayed count divided by the IO count.
% BW Host Limit Delayed IOs	The bandwidth delayed count divided by the IO count.
Host Limit Delayed Time (ms)	The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
Avg Host Limit Delayed Time (ms)	The average time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
Avg IO Host Limit Delayed Time (ms)	The IO delay time divided by the IO count.
Avg BW Host Limit Delayed Time (ms)	The bandwidth delay time divided by the bandwidth delay count.

Group metrics

This is a list of all KPIs that are available for device groups and RDF groups.

Table 18: Device group and RDF group metrics

Group metric	Descriptions
Host IOs/sec	The number of host operations performed each second by the group.
Host Reads/sec	The number of host read operations performed each second by the group.
Host Writes/sec	The number of host write operations performed each second by the group.
Host Hits/sec	The number of host read/write operations performed each second by the group that were immediately satisfied from cache.
Host Read Hits/sec	The number of host read operations performed each second by the group that were immediately satisfied from cache.
Host Write Hits/sec	The number of host write operations performed each second by the group that were immediately satisfied from cache.
Host Misses/sec	The number of host read/write operations performed each second by the group that could not be satisfied from cache.
Host Read Misses/sec	The number of host read operations performed each second by the group that were not satisfied from cache.
Host Write Misses/sec	The number of host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
Read Response Time (ms)	The average time that it took the Symmetrix to serve one read IO for this group.
Write Response Time (ms)	The average time that it took the Symmetrix to serve one write IO for this group.
Read Miss Response Time	The average time that it took the Symmetrix to serve one read miss IO for this group.

Table 18: Device group and RDF group metrics (continued)

Group metric	Descriptions
(ms)	
Write Miss Response Time (ms)	The average time that it took the Symmetrix to serve one write miss IO for this group.
RDF/S Write Response Time (ms)	A summary of the read, write, and average response times for the selected SRDF/S group.
% Reads	The percent of IO operations that were reads.
% Writes	The percent of IO operations that were writes.
% Read Hit	The percent of read operations, performed by the group, that were immediately satisfied by cache.
% Write Hit	The percent of write operations, performed by the group, that were immediately satisfied by cache.
% Read Miss	The percent of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
% Write Miss	The percent of write miss operations performed each second by the group. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the group
Seq IOs/sec	Number of IO operations performed each second that were sequential.
Seq Reads/sec	Number of read IO operations performed each second that were sequential.
Seq Writes/sec	Number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	Number of sequential read operations performed each second by the group that were immediately satisfied from cache.
Seq Read Miss/sec	Number of sequential read operations performed each second by the group that were misses.
Seq Write Hits/sec	Number of sequential write operations performed each second by the group that were immediately satisfied from cache.
Seq Write Misses/sec	Number of sequential write operations performed each second by the group that were misses.
Random IOs/sec	IOs from a host not identified as part of a sequential stream.
Random Reads/sec	Read IO commands from a host not identified as part of a sequential stream.
Random	Write IO commands from a host not identified as part of a sequential stream.

Table 18: Device group and RDF group metrics (continued)

Group metric	Descriptions
Writes/sec	
Random Read Hits/sec	Random read IOs that were satisfied from the cache.
Random Write Hits/sec	Random write IOs that were immediately placed in cache because space was available.
Random Read Misses/sec	Random read IOs that were misses.
Random Write Misses/sec	Random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec/total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec/total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec/total writes per sec)
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total ios per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total ios per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP	The maximum number of write-pending slots available for the group.

Table 18: Device group and RDF group metrics (continued)

Group metric	Descriptions
Threshold	
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes to make a request by the disk directors to the cache.
BE Disk Read Response Time (ms)	The average time it takes cache to respond to a read request by the disk directors.
BE Read Task Time (ms)	The time from the point when the HA puts the read request on the queue and the DA picks it up; this can be considered queue time.
% Hit	The percent of IO operations that were immediately satisfied from cache.
% Miss	The percent of IO operations that were misses.
% Random Reads	The percent of all read IOs that were random.
% Random Writes	The percent of all write IOs that were random.
BE % Reads	The percent of the back-end IO that were read requests.
BE % Writes	The percent of the back-end IO that were write requests.
Response Time (ms)	The average time it takes to satisfy IO requests.
BE Partial Sector Writes (KB)	The number of partial sector writes by the disk directors.
IO Density	The number of BE requests per GB of disk.
BE Optimize	These metrics are for internal Symmetrix operations.

Table 18: Device group and RDF group metrics (continued)

Group metric	Descriptions	
Writes (KB)		
BE XOR Reads		
BE XOR Read (KB)		
BE Reads for Copy		
BE Reads for Copy (KB)		
BE Reads for Copy		
BE Reads for Copy (KB)		
BE Reads for VLUN Migration		
BE Reads for VLUN Migration		
BE Reads for VLUN Migration (KB)		
BE Writes for VLUN Migration		
BE Writes for VLUN Migration (KB)		
BE Writes for Rebuild		
BE Writes for Rebuild (KB)		
BE RDF Copy		
BE RDF Copy (KB)		
BE Prefetched MBs/sec		The number of MBs prefetched from disk to cache in a second.
XtremSW Cache Read Hits/sec		The cumulative number of reads that were hits by XtremSW Cache per second.

Table 18: Device group and RDF group metrics (continued)

Group metric	Descriptions
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read for this volume.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write for this volume.
XtremSW Cache IOs/sec	The average number of IOs served by XtremSW Cache per second.
XtremSW Cache Skipped IOs/sec	The average number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
XtremSW Cache Reads	The average time it took XtremSW Cache to serve one read.
XtremSW Cache Writes	The average time it took XtremSW Cache to serve one write.
XtremSW Cache Read Hits	The number of reads served by XtremSW Cache that were hits.
XtremSW Cache MBs Read/sec	Cumulative number of host MBs read by the XtremSW Cache per second.
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache MBs/sec	Cumulative number of host MBs read and written by the XtremSW Cache per second.

Table 18: Device group and RDF group metrics (continued)

Group metric	Descriptions
XtremeSW Cache Read Response Time	The average time it took the XtremeSW Cache to serve one read for this volume.
XtremeSW Cache Write Response Time	The average time it took the XtremeSW Cache to serve one write for this volume.
XtremeSW Cache Response Time	The average time it took the XtremeSW Cache to serve one IO for this volume.
XtremeSW Cache Avg Read Size (KB)	The average size of a read served by the XtremeSW Cache.
XtremeSW Cache Avg Write Size (KB)	The average size of a write served by the XtremeSW Cache.
XtremeSW Cache Avg IO Size (KB)	The average size of an IO (read or write) served by the XtremeSW Cache.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	Number of host MBs read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: (Optimized MBs Read/sec * 1024 / Optimized Reads/sec)
Front-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to front end directors serving their IOs.

Disk Buckets Explore metrics

This is a list of all KPIs that are available for disk buckets.

Table 19: Disk Buckets metrics

New Name	Description
% Busy	The percent busy of disks in this bucket.
Avg IO/sec	The average number of disk IOs of disks in this bucket.
Avg MBs/sec	The average number of disk MBs read and written to disks in this bucket.

Table 19: Disk Buckets metrics (continued)

New Name	Description
Avg Queue Depth	As an IO enters the queue it first checks how deep the queue is. After incrementing the queue depth bucket, it increments the associated count bucket by 1. This then represents the number of IOs that encountered queues of certain depth during the past time interval. The intent of these buckets is to identify IO bursts which in turn generate large queues and long response times.
Avg Response Time	The weighted average response time (read + writes) for disks in this bucket.
Total Capacity (GB)	The total capacity of disks in this bucket.
Used Capacity (GB)	The used capacity of disks in this bucket.
Avg Number of Disks	The average number of disks in this bucket.
Reads/sec	The average number of disk Reads of disks in this bucket.
Writes/sec	The average number of disk Writes of disks in this bucket
MBs Read/sec	The average number of MBs read from disks in this bucket.
MBs Written/sec	The average number of MBs written to disks in this bucket.
Avg Read Response Time (ms)	The average read response time.
Avg Write Response Time (ms)	The average write response time.

Disk metrics

This is a list of all KPIs that are available for disks and spare disks.

Table 20: Disk/Spare disk metrics

New Name	Description
% Busy	The percent of time that the disk is busy serving IOs.
% Idle	The percent of time the disk is idle.
Avg Queue Depth	Calculated value: Accumulated queue depth/total SCSI command per sec.
Total SCSI Commands/sec	The total number of read commands, write commands, skip mask commands, verify commands, XOR write commands, and XOR write-read commands performed by the Symmetrix disk each second.
IO/sec	The number of host read and write requests for the disk.
Reads/sec	The number of host reads per second for the disk.
Writes/sec	The number of host writes per second for the disk.
MBs/sec	The size of the IO from the host to the disk per second.
MBs Read/sec	The read throughput (MBs) of the disk per second.
MBs Written/sec	The write throughput (MBs) of the disk per second.
Avg Read Size (KB)	The average number of kilobytes for a single read command.
Avg Write Size (KB)	The average number of kilobytes for a single write command.
Avg Response Time	The average response time for the reads and writes.
Avg Read Response Time (ms)	The average time it took the disk to serve one read command.
Avg Write Response Time (ms)	The average time it took the disk to serve one write command.
Seeks/sec	The number of times each second that the disk head moved to find data.
Seek Distance/sec	The number of hypervolumes that the disk head crossed (during all seeks) each second.
Avg Hypers per Seek	The average number of hypervolumes that the disk head crossed during one seek.
Verify	The number of commands that verify the integrity of the data on the disk.

Table 20: Disk/Spare disk metrics (continued)

New Name	Description
Commands/sec	
Skip Mask Commands/sec	The skip mask support offers an emulation of the ability to efficiently transfer “nearly sequential” streams of data. It allows a sequential read or write to execute but “skip over” certain unwanted or unchanged portions of the data stream, thereby transferring only those portions of the sequential stream that have changed and need to be updated. The skip mask mechanism increases throughput by saving bandwidth; both the bandwidth of processing multiple commands and the bandwidth of transferring unnecessary data.
XOR Write Commands/sec	The number of exclusive OR (XOR) write commands performed each second by the disk. XOR commands are used to establish parity protection in RAID-S and RAID 5 configurations.
XOR Read Commands/sec	The number of exclusive OR (XOR) write commands performed each second by the disk. XOR commands are used to establish parity protection in RAID-S and RAID 5 configurations.
Total Capacity (MB)	The total capacity of the disk (MBs).
Used Capacity (GB)	The total used capacity of the disk (GBs).
% Used Capacity	The percent of the disk that is used.
% Free Capacity	The percent of the disk that is free.

Disk group metrics

This is a list of all KPIs that are available for disk groups.

Table 21: Disk group metrics

Metric	Description
% Busy	The percent of time that the disk group is busy serving IOs.
% Idle	The percent of time the disk group is idle.
Total SCSI Commands/sec	The total number of read commands, write commands, skip mask commands, verify commands, XOR write commands, and XOR write-read commands performed by the disk group each second.
IOs/sec	The total number of read and write IOs per second.
Reads/sec	The number of reads per second for the disk group.
Writes/sec	The number of writes per second for the disk group.
MBs/sec	The total number of MBs per second for the disk group.
MBs Read/sec	The read throughput (MBs) of the disk group per second.
MBs Written/sec	The write throughput (MBs) of the disk group per second.
Avg Read Size (KB)	The average number of kilobytes for a single read command.
Avg Write Size (KB)	The average number of kilobytes for a single write command.
Avg Response Time	The average time it took the disk group to service IOs.
Avg Read Response Time (ms)	The average time it took the disk group to serve one read command.
Avg Write Response Time (ms)	The average time it took the disk group to serve one write command.
Total Capacity (GB)	The total capacity of all the disks in the disk group.
Used Capacity (GB)	The total capacity allocated for all the disks in the disk group.
% Used Capacity	The percent of the disk group capacity that is allocated.
% Free Capacity	The percent of the disk group capacity that is free.
IO Density	The number of BE requests per GB of disk. (BE Reads + BE Writes) /

Table 21: Disk group metrics (continued)

Metric	Description
	allocated capacity With FAST moving active extents to higher tiers, this metric is a good indication of success (the IO density on Flash tiers should be higher than the density on SATA tiers.)

Disk group storage tier metrics

This is a list of all KPIs that are available for Symmetrix standard tiers.

Table 22: Disk group storage tier metrics

Metric	Description
% Busy	The percent of time that the disk is busy serving IOs.
% Idle	The percent of time the disk is idle.
Total SCSI Commands/sec	The total number of read commands, write commands, skip mask commands, verify commands, XOR write commands, and XOR write-read commands performed by the Symmetrix disk each second.
Reads/sec	The number of reads per second for the disk.
Writes/sec	The number of writes per second for the disk.
MBs/sec	The total MBs read and written per second.
MBs Read/sec	The read throughput (MBs) of the disk per second.
MBs Written/sec	The write throughput (MBs) of the disk per second.
Avg Read Size (KB)	The average number of kilobytes for a single read command.
Avg Write Size (KB)	The average number of kilobytes for a single write command.
Avg Response Time	Calculated value: $((\text{Total Read Time/sec} + \text{Total Write Time/sec}) / (\text{Reads/sec} + \text{Writes/sec}))$
Avg Read Response Time (ms)	The average time it took the disk to serve one read command.
Avg Write Response Time (ms)	The average time it took the disk to serve one write command.
Total Capacity (GB)	The total capacity of the disk in GBs.
Used Capacity (GB)	The used capacity of the disk in GBs.
% Used Capacity	The percent of the disk capacity that is used.
% Free Capacity	The percent of the disk capacity that is free.
IO Density	The number of BE requests per GB of disk. $(\text{Reads} + \text{Writes}) / \text{total capacity}$ With FAST moving active extents to higher tiers, this metric is a good

Table 22: Disk group storage tier metrics (continued)

Metric	Description
	indication of success (the IO density on Flash tiers should be higher than the density on SATA tiers.)

FAST policy metrics

This is a list of all KPIs that are available for FAST policies.

Table 23: FAST policy metrics

Metric	Description
Allocated SG Capacity (GB)	The number of GBs of the storage group that are allocated to the FAST policy.
Total SG Capacity (GB)	The total capacity of the storage group.
Allocated Pool Capacity (GB)	The number of GBs of the virtual pool that are allocated to the FAST policy.
Total Pool Capacity (GB)	The total capacity of the virtual pool.
BE Req/sec	The number of data requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	The number of MBs transferred each second between cache and the director.
Allocated SG OOP Capacity (GB)	The GBs in the storage group that currently do not reside on the tiers defined in the FAST policy.
% Used Capacity	The percent of the pools capacity that is used.
Ingress Tracks	The number of tracks entering the pool.
Egress Tracks	The number of tracks leaving the pool.

Front-end director metrics

This is a list of all KPIs that are available for front-end directors.

Table 24: Front-end director metrics

Metric	Description
% Busy	The percent of time the directory is busy.
% Idle	The percent of time the directory is idle.
Host IOs/sec	A host command for data transfer.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Reqs/sec	Data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
Read Reqs/sec	A read data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
Write Reqs/sec	A write data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
Average Read Response Time (ms)	The calculated average response time for reads.
Average Write Response Time (ms)	The calculated average response time for writes.
Hit Reqs/sec	A request that is immediately serviced from the cache (instead of having to wait for the data to arrive, or destage from the disk).
Read Hit Reqs/sec	A read request that is immediately serviced from the cache (instead of having to wait for the data to arrive from the disk.)
Write Hit Reqs/sec	A write request that is immediately serviced from the cache (instead of having to wait for the data to be destaged to the disk.)
Miss Req/sec	A request that is a miss.
Read Miss Reqs/sec	A read request that cannot be satisfied immediately from the cache and needs to wait for the data to arrive from the disk.
Write Miss Reqs/sec	A write request that cannot be satisfied immediately from the cache and needs to wait for the data to be destaged to the disk.
% Read Reqs	The percent of read requests out of host commands for data transfer.
% Write Reqs	The percent of write requests out of host commands for data transfer.

Table 24: Front-end director metrics (continued)

Metric	Description
% Hit Reqs	The percent of requests that are served from cache.
% Read Req Hit	The percent of read requests that are served from cache.
System WP Events/sec	A write miss due to the system write pending limit having been reached.
Device WP Events/sec	A write miss due to the volume write pending limit having been reached.
Port 0 IOs/sec	The count of the IOs passed through port 0 per second.
Port 0 MBs/sec	The size of the IOs in MBs passed through port 0 per second.
Port 0 Avg IO Size (KB)	The average IO size served by port 0.
Port 1 IOs/sec	The count of the IOs passed through port 1 per second.
Port 1 MBs/sec	The size of the IOs in MBs passed through port 1 per second.
Port 1 Avg IO Size (KB)	The average IO size served by port 1.
Syscall Count/sec	The total number of calls seen by this director.
Avg Time per Syscall	The average time that it took to serve one system call.
Syscall Remote Dir Count/sec	The number of calls per second sent from the local director to another director in the same system.
Syscall RDF Dir count/sec	The number of calls per second sent via RDF to a remote system.
Slot Collisions/sec	The number of write misses due to the desired cache slot being locked by another request.
Queue Depth Count Range 0 - Queue Depth Count Range 9	As an IO enters the queue it first checks how deep the queue is. After incrementing the queue depth bucket, it increments the associated count bucket by 1. This then represents the number of IOs that encountered queues of certain depth during the past time interval. The intent of these buckets is to identify IO bursts which in turn generate large queues and long response times.
Avg Queue Depth Range	As an IO enters the queue it first checks how deep the queue is. Based on depth, the applicable queue depth bucket increments with the value seen by

Table 24: Front-end director metrics (continued)

Metric	Description
0 - Avg Queue Depth Range 9	the IO. For example, an IO that encounters a queue depth of 7 will increment bucket #2 (depth 5-9 for OS or 7-14 for MF) by 7. The intent of these buckets is to identify IO bursts which in turn generate large queues and long response times.
% Write Req Hit	The percent of write requests that were satisfied from cache.
Read RT Count Range 0 to Read RT Count Range 64	These buckets show the distribution of the number of reads to the FE directors over the specified time range.
Acc Read RT Count Range 32 to 64	The accumulated number of reads from 32 to 64.
Write RT Count Range 0 to Write RT Count Range over 64	These buckets show the distribution of the number of writes to the FE directors over the specified time range.
Total Read Count	The total number of reads to the FE directors.
Total Write Count	The total number of writes to the FE directors.
Port 0 MBs Read/sec	The number of MBs per second on port 0 that were reads.
Port 1 MBs Read/sec	The number of MBs per second on port 1 that were reads.
Port 0 MBs Written/sec	The number of MBs per second on port 0 that were writes.
Port 1 MBs Written/sec	The number of MBs per second on port 1 that were writes.
% Busy Port 0	The utilization percent for port 0.
% Busy Port 1	The utilization percent for port 1.
Port 0 Speed Gb/sec	The number of gigabits moving through port 0 per second.
Port 1 Speed Gb/sec	The number of gigabits moving through port 1 per second.

Table 24: Front-end director metrics (continued)

Metric	Description
Average Read Response Time 0 to 1 to Average Read Response Time over 64	These buckets show the distribution of the total average read response time.
Average Write Response Time 0 to 1 to Average Write Response Time over 64	These buckets show the distribution of the total average write response time.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	The number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: $(\text{Optimized MBs Read/sec} * 1024 / \text{Optimized Reads/sec})$
Port 0 Reads/sec	The number of read operations through port 0 per second.
Port 1 Reads/sec	The number of read operations through port 1 per second.
Port 0 Writes/sec	The number of write operations through port 0 per second.
Port 1 Writes/sec	The number of write operations through port 1 per second.
Host IO Limit IOs/sec	The number of IOs per seconds that are executed by the director under Host IO quota limits.
Host IO Limit MBs/sec	The throughput in MBs that are transferred by the director under Host IO quota limits.
Host IO Limit % Quota Delayed IOs	The percent of IOs delayed due to Host IO quota limits.
Host IO Limit Delayed Time (ms)	The time in milliseconds that IOs are delayed because Host IO quota limits.

Front-end director Real Time metrics

This is a list of all KPIs that are available for front-end directors in Real Time.

Table 25: Front-end director Real Time metrics

Metric	Description
% Busy	The percent of time the directory is busy.
Host IOs/sec	The number of host commands for data transfer per second.
Host MBs/sec	The size (MBs) of the host throughput per second.
Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
Read Reqs/sec	A read data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
Write Reqs/sec	A write data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment or both. The requests rate should be either equal to or greater than the IO rate.
Avg Read Response Time (ms)	The calculated average response time for reads.
Avg Write Response Time (ms)	The calculated average response time for writes.
Hits Reqs/sec	A request that is immediately serviced from the cache (instead of having to wait for the data to arrive, or destage, from the disk.)
Read Hit Reqs/sec	A read request that is immediately serviced from the cache (instead of having to wait for the data to arrive from the disk.)
Write Hit Reqs/sec	A write request that is immediately serviced from the cache (instead of having to wait for the data to be destaged to the disk.)
Read Miss Reqs/sec	A read request that cannot be satisfied immediately from the cache and needs to wait for the data to arrive from the disk.
Write Miss Reqs/sec	A write request that cannot be satisfied immediately from the cache and needs to wait for the data to destage to the disk.
System WP Events/sec	A write miss due to the system write pending limit having been reached.
Device WP Events/sec	A write miss due to the volume write pending limit having been reached.
Port 0 IOs/sec	The number of IOs that pass through port 0 per second.

Table 25: Front-end director Real Time metrics (continued)

Metric	Description
Port 0 MBs/sec	The size of the IOs in MBs that pass through port 0 per second.
Port 1 IOs/sec	The number of IOs that pass through port 1 per second.
Port 1 MBs/sec	The size of the IOs in MBs that pass through port 1 per second.
Acc Director Idle Time	The accumulated time the director was idle.
Read RT Count Range 0 to Read RT Count Range 64	These buckets show the distribution of the number of reads to the FE directors over the specified time range.
Write RT Count Range 0 to Write RT Count Range 64	These buckets show the distribution of the number of writes to the FE directors over the specified time range.
Total Read Count	The total number of read requests to the FE directors.
Total Write Count	The total number of write requests to the FE directors.
% Busy Port 0	The percent of time Port 0 was busy.
% Busy Port 1	The percent of time Port 1 was busy.
Avg Read Response Time 0 to Avg Read Response Time 64	A calculated average response time for reads.
Avg Write Response Time 0 to Avg Write Response Time 64	A calculated average response time for writes.
Avg Read Response Time Over 64	A calculated average response time for reads.
Avg Write Response Time Over 64	A calculated average response time for writes.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	Number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: (Optimized MBs Read/sec * 1024 / Optimized Reads/sec)

Table 25: Front-end director Real Time metrics (continued)

Metric	Description
Port 0 Reads/sec	The number of read operations through port 0 per second.
Port 1 Reads/sec	The number of read operations through port 1 per second.
Port 0 Writes/sec	The number of write operations through port 0 per second.
Port 1 Writes/sec	The number of write operations through port 1 per second.

Front-end port metrics

This is a list of all KPIs that are available for front-end ports that are associated with a storage group with Host IO Limits set on the directors.

Table 26: Front-end ports metrics

Metric	Description
Host IOs/sec	The number of IOs the port is performing on behalf of the specific SG.
IO Limit IOs/sec	The sum of the IOs that were delayed due to the SG's IO limits.
BW Limit Delayed IOs/sec	The sum of the MBs that were delayed due to the SG's bandwidth limits.
Host MBs/sec	The number of MBs the port is processing on behalf of the specific SG.
Host Limit Delayed IOs/sec	The sum of the IOs that were delayed due to the SG's IO limits and bandwidth limits.
% Host Limit Delayed IOs	The host limit delayed IOs divided by the IO count.
% IO Host Limit Delayed IOs	The IO delayed count divided by the IO count.
% BW Host Limit Delayed IOs	The bandwidth delayed count divided by the IO count.
Host Limit Delayed Time (ms)	The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
Avg Host Limit Delayed Time (ms)	The average time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
Avg IO Host Limit Delayed Time (ms)	The IO delay time divided by the IO count.
Avg BW Host Limit Delayed Time (ms)	The bandwidth delay time divided by the bandwidth delay count.

FE Port SG metrics

This is a list of all KPIs that are available for Host IO limits set on front-end ports for the storage group.

Table 27: FE Port by SG metrics

Metric	Description
Host IO Limit IOs/sec	The number of IOs the port is performing on behalf of the specific SG.
Host IO Limit MBs/sec	The number of MBs the port is processing on behalf of the specific SG.
Host IO Limit Delayed IOs/sec	The time the IOs are delayed because of the IO limit set for the SG.
Host IO Limit Delayed Time (ms)	The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
Host IO Limit % Delayed IOs	The quota delayed IOs divided by the IO count.
Host IO Limit Avg Delayed Time	The average time that IOs are delayed (due to the SG's set IO and bandwidth quotas) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth quotas).
IOs/sec	The number of IOs per second served.
% Busy	The percent of time the port is busy.
MBs/sec	The throughput for the port.

RDF director metrics

This is a list of all KPIs that are available for RDF directors.

Table 28: RDF director metrics

Metric	Description
% Busy	The percent of time the directory is busy.
% Idle	The percent of time the directory is idle.
IO/sec	The average IO count per second for the director.
Avg IO Service Time (ms)	The average time the director takes to serve IO.
Reqs/sec	The total read and write requests per second.
Read Reqs/sec	The read requests per second.
Write Reqs/sec	The write requests per second.
MBs Sent and Recieved/sec	The total read and write MBs on the RDF director per second.
MBs Sent/sec	The size of the host data transfer in MBs for the director.
MBs Received/sec	The size of the data received in MBs for the director.
% Busy Port 0	The percent of time that port 0 is busy.
Avg IO Size Received (KB)	Calculated value: (MBs received per sec / reads per sec)
Avg IO Size Sent (KB)	Calculated value: (MBs sent per sec / writes per sec)
Syscall Count/sec	The total number of calls seen by this director.
Syscall Time/sec	The time to satisfy the calls by this director.
Syscall Remote Dir Count/sec	The number of calls sent from the local director to another director in the same system.
Syscall RDF Dir Count/sec	The number of calls sent via RDF to a remote system.
Avg Time per Syscall	Calculated value: (syscall time / syscall count)
Write Miss Reqs/sec	The write misses for the RDF director per second. A miss occurs when the request has to wait for data to destage from cache to disk.
Port 0 Speed Gb/sec	The number of Gbits per second through port 0.
Tracks Received/sec	The number of tracks received by this director per second.
Tracks Sent/sec	The number of tracks sent by this director per second.
Number of Links	The number of links managed by this RDF director.

Table 28: RDF director metrics (continued)

Metric	Description
Number of Compressed Links	The number of compressed links on this RDF director.
Compressed MBs Received/sec	The number of compressed MBs/sec received via this RDF director.
Compressed MBs Sent/sec	The number of compressed MBs/sec sent via this RDF director.
Compressed MBs Sent and Received/sec	The number of compressed MBs/sec sent and received via this RDF director.
% Compressed MBs Sent and Received	The percent of MBs sent and received that are compressed.
% Compressed MBs Received	The percent of MBs received from R2 that are compressed.
% Compressed MBs Sent	The percent of MBs sent to R2 that are compressed.
Copy IOs/sec	The number of copy IOs per second via this RDF director.
Copy MBs/sec	The number of copy MBs per second via this RDF director.
Sync Write Reqs/sec	The number of sync write requests managed by this RDF director per second.
Sync MBs Sent/sec	The number of sync MBs sent per second via this RDF director.
Async Write Reqs/sec	The number of async write requests managed by this RDF director per second.
Async MBs Sent/sec	The number of async MBs sent per second via this RDF director.

RDF director Real Time metrics

This is a list of all KPIs that are available for RDF directors in Real Time.

Table 29: RDF director Real Time metrics

Metric	Description
% Busy	The percent of time that a director is busy.
IOs/sec	The number of IO commands to the disk per second.
Write Reqs/sec	A data transfer of a write between the cache and the director.
Tracks Received/sec	The number of tracks that received data each second.
Tracks Sent/sec	The number of tracks that sent data each second.

RDF/A group metrics

This is a list of all KPIs that are available for RDF/A groups.

Table 30: RDF/A group metrics

Metric	Description
Avg Cycle Time	The mean time (in seconds) of the last 16 cycles. (Calculated as true averages- e.g. if only 8 cycles have occurred, the average represents those 8 cycles only.)
Avg Cycle Size	The average size of all the cycles in RDF/A sessions.
Duration of Last Cycle	The cycle time (in seconds) of the most recently completed cycle. It should be noted that in a regular case the cycles switch every ~30 sec, however, in most cases the collection interval is in minutes, which means some cycle times will be skipped. This an important counter to look at to figure out if SRDF/A is working as expected.
Total HA Writes/sec	The total host adapter writes, measured in write commands to SRDF/A volumes only.
Total HA Repeat Writes/sec	The total number of writes to a slot already in the active cycle. Total host adapter repeat writes, measured in write commands to SRDF/A volumes only. This counter helps estimate the cache locality of reference, i.e., how much cache is saved by the re-writes. This does not give any indication to the bandwidth locality of reference.
RDF R1 to R2 IOs/sec	The number of IOs/sec being transmitted across the links for this session.
RDF R1 to R2 MBs/sec	The MB/sec being transmitted across the links for this session.
RDF R2 to R1 IOs/sec	The number of IOs/sec being transmitted across the links for this session.
RDF R2 to R1 MBs/sec	The MB/sec being transmitted across the links for this session.
Host Writes/sec	The number of write IOs/sec coming in from the hosts for the volumes in this session.
Host MBs Written/sec	The throughput in MB/sec of writes coming in from the hosts for the volumes in this session.
System WP Limit	The maximum number of write-pending slots for the entire Symmetrix system.
System WP Count	The number of cache slots marked as write pending for the system.
Local WP Count	The number of write pending slots waiting to be de-staged to disk. On the R1 - should be less than or equal to the system write pending count.

Table 30: RDF/A group metrics (continued)

Metric	Description
RDF/A WP Count	The number of RDF/A writes pending.
% RDF/A Max Cache Used	A parameter establishing the percent of system write pendings at which an RDF/A session will drop.
Min Cycle Time	The setting for the minimum number of seconds for a cycle.
DSE Used Tracks	The active cycle spilled count plus the inactive cycle spilled count.
% R1 Used Cache	The ratio of system write pending count to the maximum write pending slots. This is the percent of cache used by this RDF/A group.
% R2 Used Cache	The ratio of system write pending count to the maximum write pending slots. This is the percent of cache used by this RDF/A group.
Reads/sec	The number of tracks read from the R2 (host reads + copy).
MBs Read/sec	The total amount of data read from the R2 in MBs+G5.
Avg IO Service Time (ms)	The average time it took to satisfy reads and writes.
Time Since Last Switch	The time (in seconds) since the last switch from active to inactive or vice versa. Since in a regular case the cycles switch every ~30 sec and the samples are taken every few minutes, this counter may not show very significant data, however, when cycles elongate beyond the sample time, this counter can help indicate an estimate of the R1 to R1 lag time.
Active Cycle Size	The count of cache slots that measures the amount of data captured (on the R1 side) or received (on the R2 side).
Inactive Cycle Size	The count of cache slots that measures the amount of data transmitted (on the R1 side) or applied (on the R2 side).
RDF/A Session Indicator	Indicates if group is RDF/A (0/1).
RDF/A Active Indicator	Indicates if RDF/A group is active (0/1).
Session Priority	The order in which SRDF/A groups are dropped out of SRDF/A in the event of a cache limit being reached. The highest priority is 1, the lowest is 64, and the default is 33.
Cycle	The active cycle number identifier for the given SRDF/A session. In a regular

Table 30: RDF/A group metrics (continued)

Metric	Description
Number	case, the cycles switch every ~30 seconds, however, in most cases the collection interval is in minutes. Calculating the difference in cycle numbers between collection intervals relative to the minimum cycle time will show how many cycles occurred over the last interval.
High Water Mark	The greatest number of cache slots transmitted (on the R1 side) or applied (on the R2 side) in any cycle so far in the session.
R1 to R2 Lag Time	The time that R2 is behind R1 (RPO). This is calculated as the last cycle time plus the time since last switch. In a regular case, the cycles switch every ~30 sec and the samples are taken every few minutes, therefore this counter may not show very significant data, however, when cycles elongate beyond the sample time, this counter can help indicate an estimate of the RPO.
Num Devs in Session	The number of Symmetrix volumes defined in this session.
R2 Last Restore Time	The time, in seconds, that the restore of the last active cycle took to complete. The restore is an operation done by the DA and RA to mark the tracks in the apply cycle as Write Pending to the local drives. This operation is usually a very quick cache operation, unless one of the volumes reaches its WP limit. The importance of this counter is that it shows one of the common reasons for an increase in the cycle time.
R2 Avg Restore Time	The mean time (in secs) of the last 16 restores. (Calculated as true averages, e.g., if only 8 restores have occurred, the average represents those 8 restores only.)
Consistency Protection	Specifies whether consistency protection is set.
DSE Threshold	The percent of write pendings before DSE activates.
Writes/sec	The number of write requests to the RDF director per second.
Session Number	The number of the SRDF/A session.
Uncommitted Tracks	The number of tracks for all the SRDF/A volumes in the session that have not been committed to the R2 volume.
MBs Written/sec	The number of MBs written per second for the SRDF/A session.
Read Hits/sec	The total number of read operations that were immediately satisfied by cache.
Max Cycle Time	The setting for the maximum number of seconds for a cycle.
Number of	The number of links used by this RDF/A group.

Table 30: RDF/A group metrics (continued)

Metric	Description
Links	
Number of Compressed Links	The number of compressed links used by this RDF/A group.
Compressed MBs Received/sec	The number of compressed MBs/sec received by this RDF/A group.
Compressed MBs Sent/sec	The percent of compressed MBs/sec sent by this RDF/A group.
Compressed MBs Sent and Received/sec	The number of compressed MBs sent and received by this RDF/A group.
% Compressed MBs Received	Calculated value: $((\text{Compressed MBs Received/sec} / \text{MBs Received/sec}) * 100)$
% Compressed MBs Sent	Calculated value: $((\text{Compressed MBs Sent/sec} / \text{MBs Sent/sec}) * 100)$
% Compressed MBs Sent and Received	Calculated value: $((\text{Compressed MBs Sent and Received/sec} / \text{MBs Sent and Received/sec}) * 100)$

Snap and DSE pool metrics

This is a list of all KPIs that are available for Snap and DSE pools.

Table 31: Snap and DSE pool metrics

Metric	Description
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Read Request Time (ms)	The average time it takes to make a read request by the disk directors to the cache.
BE Disk Read Response Time (ms)	The average time it takes cache to respond to a read request by the disk directors.
BE Read Task Time (ms)	The time from the point when the HA puts the read request on the queue and the DA picks it up; this can be considered queue time.
% Used Capacity	The percent of the pools total capacity that is used.
Total Pool Capacity (GB)	The total pool capacity in GBs.
Enabled Pool Capacity (MB)	The enabled pool capacity in GBs.
Used Pool Capacity (GB)	The used pool capacity in GBs.
Allocated Pool Capacity (MB)	The allocated pool capacity in GBs.

Storage group explore metrics

This is a list of all KPIs that are available for storage groups.

Table 32: Storage group explore metrics

Group metric	Descriptions
Host IOs/sec	The number of host operations performed each second by the group.
Host Reads/sec	The number of host read operations performed each second by the group.
Host Writes/sec	The number of host write operations performed each second by the group.
Host Hits/sec	The number of host read/write operations performed each second by the group that were immediately satisfied from cache.
Host Read Hits/sec	The number of host read operations performed each second by the group that were immediately satisfied from cache.
Host Write Hits/sec	The number of host write operations performed each second by the group that were immediately satisfied from cache.
Host Misses/sec	The number of host read/write operations performed each second by the group that could not be satisfied from cache.
Host Read Misses/sec	The number of host read operations performed each second by the group that were not satisfied from cache.
Host Write Misses/sec	The number of host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
Read Response Time (ms)	The average time that it took the Symmetrix to serve one read IO for this group.
Write Response Time (ms)	The average time that it took the Symmetrix to serve one write IO for this group.
Read Miss Response Time (ms)	The average time that it took the Symmetrix to serve one read miss IO for this group.
Write Miss Response Time	The average time that it took the Symmetrix to serve one write miss IO for this group.

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
(ms)	
RDF/S Write Response Time (ms)	A summary of the read, write, and average response times for the selected SRDF/S group.
% Reads	The percent of IO operations that were reads.
% Writes	The percent of IO operations that were writes.
% Read Hit	The percent of read operations, performed by the group, that were immediately satisfied by cache.
% Write Hit	The percent of write operations, performed by the group, that were immediately satisfied by cache.
% Read Miss	The percent of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
% Write Miss	The percent of write miss operations performed each second by the group. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the group
Seq IOs/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the group that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the group that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the group that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the group that were misses.
Random IOs/sec	IOs from a host not identified as part of a sequential stream.
Random Reads/sec	Read IO commands from a host not identified as part of a sequential stream.
Random Writes/sec	Write IO commands from a host not identified as part of a sequential stream.

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
Random Read Hits/sec	Random read IOs that were satisfied from the cache.
Random Write Hits/sec	Random write IOs that were immediately placed in cache because space was available.
Random Read Misses/sec	Random read IOs that were misses.
Random Write Misses/sec	Random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	The percent of IO operations that were sequential.
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
Max WP Threshold	The maximum number of write-pending slots available for the group.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	Total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes to make a request by the disk directors to the cache.
BE Disk Read Response Time (ms)	The average time it takes cache to respond to a read request by the disk directors.
BE Read Task Time (ms)	The time from the point when the HA puts the read request on the queue and the DA picks it up; this can be considered queue time.
% Hit	The percent of IO operations that were immediately satisfied from cache.
% Miss	The percent of IO operations that were misses.
% Random Reads	The percent of all read IOs that were random.
% Random Writes	The percent of all write IOs that were random.
BE % Reads	The percent of the back-end IO that were read requests.
BE % Writes	The percent of the back-end IO that were write requests.
BE Partial Sector Writes (KB)	The KBs of back-end IO that were partial sector writes.
IO Density	The number of BE requests per GB of disk.
BE Optimized Writes (KB)	These metrics are for internal Symmetrix operations.

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
BE XOR Reads	
BE XOR Reads (KB)	
BE Reads for Copy	
BE Reads for Copy (KB)	
BE Writes for Copy	
BE Writes for Copy (KB)	
BE Reads for VLUN Migration	
BE Reads for VLUN Migration (KB)	
BE Writes for VLUN Migration	
BE Writes for VLUN Migration (KB)	
BE Writes for Rebuild	
BE Writes for Rebuild (KB)	
BE RDF Copy	
BE RDF Copy (KB)	
Device Capacity	
Device Block Size	
BE Prefetched MBs/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
Total Tracks	The total number of tracks for the storage group.
Compressed	The number of compressed tracks for the storage group.

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
Tracks	
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read for this volume.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write for this volume.
XtremSW Cache IOs/sec	The average number of IOs served by XtremSW Cache per second.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
% XtremSW Cache Reads	The number of read operations served by XtremSW Cache per second.
% XtremSW Cache Writes	The number of write operations served by XtremSW Cache per second.
% XtremSW Cache Read Hits	The number of read operations that were hits served by XtremSW Cache per second.
XtremSW Cache MBs	Cumulative number of host MBs read by the XtremSW Cache per second.

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
Read/sec	
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache MBs/sec	Cumulative number of host MBs read and written by the XtremSW Cache per second.
XtremSW Cache Read Response Time	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write Response Time	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache Response Time	The average time it took the XtremSW Cache to serve one IO.
XtremSW Cache Avg Read Size (KB)	The average size of reads (in KBs) handled by XtremSW Cache.
XtremSW Cache Avg Write Size (KB)	The average size of writes (in KBs) handled by XtremSW Cache.
XtremSW Cache Avg IO Size (KB)	The average size of IO (in KBs) handled by XtremSW Cache.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	Number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: (Optimized MBs Read/sec * 1024 / Optimized Reads/sec)
Workload Score	A combined score metric that reflects the overall contribution of a volume or list of volumes to the array.
Front-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to front-end directors serving their IOs.
Back-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to back-end directors serving their IOs.
Disk Workload	A score metric that reflects the contribution of a volume or list of volumes to

Table 32: Storage group explore metrics (continued)

Group metric	Descriptions
Score	disks serving their IOs.

Symmetrix system Historical metrics

This is a list of all KPIs that are available for Symmetrix systems.

Table 33: Symmetrix system Historical metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second by all Symmetrix volumes, including writes and random and sequential reads.
Host Reads/sec	The number of host read operations performed each second by all Symmetrix volumes.
Host Writes/sec	The number of host write operations performed each second by all Symmetrix volumes.
Host MBs/sec	The number of host MBs written and read by all of the Symmetrix volumes each second.
Host MBs Read/sec	The number of host MBs read by all of the Symmetrix volumes each second.
Host MBs Written/sec	The number of host MBs written by all of the Symmetrix volumes each second.
FE Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
Avg Read Response Time (ms)	The calculated system read response time.
Avg Write Response Time (ms)	The calculated system write response time.
% Reads	The percent of total read IO operations performed each second by all of the Symmetrix volumes.
% Writes	The percent of total write IO operations performed by all of the Symmetrix volumes.
% Hit	The percent of IO operations performed by all of the Symmetrix volumes, for which the read data was in cache and the write operation could be sent directly to cache without having to wait for data to be destaged from cache to the disks.
FE Read Reqs/sec	A read data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
FE Write Reqs/sec	A write data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.

Table 33: Symmetrix system Historical metrics (continued)

Metric	Description
BE IOs/sec	The total IO from all BE directors to the disks per second.
BE Reqs/sec	A data transfer of a read or write between the cache and the director.
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
BE Write Reqs/sec	A data transfer of a write between the cache and the director.
System WP Events/sec	The number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific Symmetrix volume was reached.
WP Count	The number of system cache slots that are write pending.
System Max WP limit	The percent of the target % at which writes are delayed. The range is from 40% to 80%.
% Cache WP	The percent of system cache that is write pending.
Avg Fall Thru Time	the average time it takes a cache slot in LRU0 to be freed up. It is the average time from the first use of the contents to its reuse by another address.
FE Hit Req/sec	The total requests from all FE directors per second that were satisfied from cache.
FE Read Hit Reqs/sec	The total read requests from all FE directors per second that were satisfied from cache.
FE Write Hit Reqs/sec	The total write requests from all FE directors per second that were satisfied from cache.
Prefetched Tracks/sec	The tracks per second pre-fetched from disk to cache upon detection of a sequential read stream.
Destaged Tracks/sec	The tracks per second saved into disks.
FE Read Miss Reqs/sec	The total read requests from all FE directors per second that were misses. A miss occurs when the requested read data is not found in cache.
FE Write Miss Reqs/sec	The total write requests from all FE directors per second that were misses. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.

Symmetrix system metrics

This is a list of all KPIs that are available for Symmetrix systems.

Table 34: Symmetrix system metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second by all Symmetrix volumes, including writes and random and sequential reads.
Host Reads/sec	The number of host read operations performed each second by all Symmetrix volumes.
Host Writes/sec	The number of host write operations performed each second by all Symmetrix volumes.
Host MBs/sec	The number of host MBs written and read by all of the Symmetrix volumes each second.
Host MBs Read/sec	The number of host MBs read by all of the Symmetrix volumes each second.
Host MBs Written/sec	The number of host MBs written by all of the Symmetrix volumes each second.
FE Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
Avg Read Response Time (ms)	The calculated system read response time.
Avg Write Response Time (ms)	The calculated system write response time.
% Reads	The percent of total read IO operations performed each second by all of the Symmetrix volumes.
% Writes	The percent of total write IO operations performed by all of the Symmetrix volumes.
% Hit	The percent of IO operations performed by all of the Symmetrix volumes, for which the read data was in cache and the write operation could be sent directly to cache without having to wait for data to be destaged from cache to the disks.
FE Read Reqs/sec	A read data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
FE Write Reqs/sec	A write data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.

Table 34: Symmetrix system metrics (continued)

Metric	Description
BE IOs/sec	The total IO from all BE directors to the disks per second.
BE Reqs/sec	A data transfer of a read or write between the cache and the director.
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
BE Write Reqs/sec	A data transfer of a write between the cache and the director.
System WP Events/sec	The number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific Symmetrix volume was reached.
WP Count	The number of system cache slots that are write pending.
System Max WP limit	The percent of the target % at which writes are delayed. The range is from 40% to 80%.
% Cache WP	The percent of system cache that is write pending.
Avg Fall Thru Time	The average time it takes a cache slot in LRU0 to be freed up. It is the average time from the first use of the contents to its reuse by another address.
FE Hit Req/sec	The number of total requests from all FE directors per second that were satisfied from cache.
FE Read Hit Reqs/sec	The number of total read requests from all FE directors per second that were satisfied from cache.
FE Write Hit Reqs/sec	The number of total write requests from all FE directors per second that were satisfied from cache.
Prefetched Tracks/sec	The number of tracks per second pre-fetched from disk to cache upon detection of a sequential read stream.
Destaged Tracks/sec	The number of tracks per second saved into disks.
FE Read Miss Reqs/sec	The number of total read requests from all FE directors per second that were misses. A miss occurs when the requested read data is not found in cache.
FE Write Miss Reqs/sec	The number of total write requests from all FE directors per second that were misses. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.

Table 34: Symmetrix system metrics (continued)

Metric	Description
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read for this volume.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write for this volume.
XtremSW Cache IOs/sec	The average number of IOs served by XtremSW Cache per second.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
XtremSW Cache Reads	The number of read operations served by XtremSW Cache per second.
XtremSW Cache Writes	The number of write operations served by XtremSW Cache per second.

Table 34: Symmetrix system metrics (continued)

Metric	Description
XtremSW Cache Read Hits	The number of read operations served by XtremSW Cache per second that were hits.
XtremSW Cache MBs Read/sec	Cumulative number of host MBs read by the XtremSW Cache per second.
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache Read Time	The average time it took the XtremSW Cache to serve one read for this volume.
XtremSW Cache Write Time	The average time it took the XtremSW Cache to serve one read for this volume.

Symmetrix system Real Time metrics

Table 35: Symmetrix system Real Time metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second by all Symmetrix volumes, including writes and random and sequential reads.
Host MBs/sec	The number of host MBs written and read by all of the Symmetrix volumes each second.
FE Reqs/sec	A data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
Avg Read Response Time (ms)	A calculated average response time for reads.
Avg Write Response Time (ms)	A calculated average response time for writes.
FE Read Reqs/sec	A read data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
FE Write Reqs/sec	A write data transfer between the director and the cache. An IO may require multiple requests depending on IO size, alignment, or both. The requests rate should be either equal to or greater than the IO rate.
BE IOs/sec	The total IO from all BE directors to the disks per second.
BE Reqs/sec	A data transfer of a read or write between the cache and the director.
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
BE Write Reqs/sec	A data transfer of a write between the cache and the director.
System WP Events/sec	The number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific Symmetrix volume was reached.
WP Count	The number of system cache slots that are write pending.
System Max WP Limit	The maximum write pending limit for the Symmetrix system.

Table 35: Symmetrix system Real Time metrics (continued)

Metric	Description
% Cache WP	The percent of system cache that is write pending.
FE Hit Reqs/sec	The total requests from all FE directors per second that were satisfied from cache.
FE Read Hit Reqs/sec	The total read requests from all FE directors per second that were satisfied from cache.
FE Write Hit Reqs/sec	The total write requests from all FE directors per second that were satisfied from cache.
FE Read Miss Reqs/sec	The total read requests from all FE directors per second that were misses.
FE Write Miss Reqs/sec	The total write requests from all FE directors per second that were misses.

Standard tier metrics

This is a list of all KPIs that are available for Symmetrix standard tiers.

Table 36: Symmetrix Tier metrics

Metric	Description
% Busy	The percent of time that the disk is busy serving IOs.
% Idle	The percent of time the disk is idle.
Avg Queue Depth	Calculated value: (Accumulated queue depth / total SCSI command per sec)
Total SCSI Commands/sec	The total number of read commands, write commands, skip mask commands, verify commands, XOR write commands, and XOR write-read commands performed by the Symmetrix disk each second.
Reads/sec	The number of reads per second for the disk.
Writes/sec	The number of writes per second for the disk.
MBs Read/sec	The read throughput (MBs) of the disk per second.
MBs Written/sec	The write throughput (MBs) of the disk per second.
Avg Read Size (KB)	The average number of kilobytes for a single read command.
Avg Write Size (KB)	The average number of kilobytes for a single write command.
Avg Read Time (ms)	The average time it took the disk to serve one read command.
Avg Write Time (ms)	The average time it took the disk to serve one write command.
Seeks/sec	The number of times each second that the disk head moved to find data.
Seek Distance/sec	The number of hypervolumes that the disk head crossed (during all seeks) each second.
Avg Hypers per Seek	The average number of hypervolumes that the disk head crossed during one seek.
Verify Commands/sec	The number of commands that verify the integrity of the data on the disk.
Skip Mask Commands/sec	The skip mask support offers an emulation of the ability to efficiently transfer “nearly sequential” streams of data. It allows a sequential read or write to execute but “skip over” certain unwanted or unchanged portions of the data stream, thereby transferring only those portions of the sequential stream that have changed and need to be updated. The skip mask mechanism increases throughput by saving bandwidth; both the bandwidth of processing

Table 36: Symmetrix Tier metrics (continued)

Metric	Description
	multiple commands and the bandwidth of transferring unnecessary data.
XOR Write Commands/sec	The number of exclusive OR (XOR) write commands performed each second by the disk. XOR commands are used to establish parity protection in RAID-S and RAID 5 configurations.
XOR Read Commands/sec	The number of exclusive OR (XOR) write commands performed each second by the disk. XOR commands are used to establish parity protection in RAID-S and RAID 5 configurations.
Total Capacity (GB)	The total capacity of the tier in GBs.
Used Capacity (GB)	The used capacity of the tier in GBs.
% Used Capacity	The percent of the tier capacity that is used.
% Free Capacity	The percent of the tier capacity that is free.

Tier by Storage Group metrics

This is a list of all KPIs that are available for Symmetrix virtual pool tiers.

Table 37: Tier by Storage Group metrics

Metric	Description
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
Allocated Capacity	The total allocated SG capacity.
IO Density	The number of BE requests per GB of disk. (BE Reads + BE Writes) / allocated capacity With FAST moving active extents to higher tiers, this metric is a good indication of success (the IO density on Flash tiers should be higher than the density on SATA tiers.)
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	The MBs transferred per second to the disk directors.
BE MBs Read/sec	The number of MBs read each second by the disk directors.

Table 37: Tier by Storage Group metrics (continued)

Metric	Description
BE MBs Written/sec	The number of MBs written each second by the disk directors.
Prefetched Tracks/sec	The number of tracks prefetched from disk to cache in a second.
Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Total Tracks	The number of tracks in the storage group.
Compressed Tracks	The number of compressed tracks in the storage group.
% Compressed Tracks	The percent of the total tracks that are compressed.

Thin pool metrics

This is a list of all KPIs that are available for thin pools.

Table 38: Thin pool metrics

Metric	Description
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	The MBs transferred per second to the disk directors.
BE MBs Read/sec	The number of MBs read each second by the disk directors.
BE MBs Written/sec	The number of MBs written each second by the disk directors.
Read Miss Response Time (ms)	The average time that it took the Symmetrix system to serve one read miss IO for this pool.
Write Miss Response Time (ms)	The average time that it took the Symmetrix system to serve one write miss IO for this pool.
RDF/S Write Response Time (ms)	The average time it took the Symmetrix system to respond to the write request from the RDF/S group.
BE Read Request Time (ms)	The average time it takes to make a read request by the disk directors to the cache.
BE Disk Read Response Time (ms)	The average time it takes cache to respond to a read request by the disk directors.
BE Read Task Time (ms)	Time from the point when the HA puts the read request on the queue and the DA picks it up - can be considered queue time.
% Used Capacity	The percent of the pools capacity that is used.
Total Pool Capacity (GB)	The total pool capacity in GBs.
Enabled Pool	The enabled pool capacity in GBs.

Table 38: Thin pool metrics (continued)

Metric	Description
Capacity (GB)	
Used Pool Capacity (GB)	The used pool capacity in GBs.
Allocated Pool Capacity (GB)	The allocated pool capacity in GBs.
BE Read Response Time (ms)	A calculated value of the response time for all back-end read requests.
BE Write Response Time (ms)	A calculated value of the response time for all back-end write requests.
BE Avg Response Time (ms)	A calculated value of the response time for all back-end read and write requests.
BE % Reads	The percent of the back-end requests that were read requests.
BE % Writes	The percent of the back-end requests that were write requests.
Ingress Tracks	The number of tracks entering the pool.
Egress Tracks	The number of tracks leaving the pool.
IO Density	The number of BE requests per GB of disk. (BE Reads + BE Writes) / allocated capacity With FAST moving active extents to higher tiers, this metric is a good indication of success (the IO density on Flash tiers should be higher than the density on SATA tiers.)
Total Tracks	The number of tracks in the thin pool.
Compressed Tracks	The number of compressed tracks in the thin pool.
% Compressed Tracks	The percent of the total tracks that are compressed.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read.
XtremSW Cache Writes /sec	The average time it took XtremSW Cache to serve one write.

Table 38: Thin pool metrics (continued)

Metric	Description
XtremSW Cache IOs/sec	The average time it took XtremSW Cache to serve one IO.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
% XtremSW Cache Reads	The percent of XtremSW Cache IO that were reads.
% XtremSW Cache Writes	The percent of XtremSW Cache IO that were writes.
% XtremSW Cache Read Hits	The percent of XtremSW Cache IO that were read hits.
XtremSW Cache MBs Read/sec	Cumulative number of host MBs read by the XtremSW Cache per second.
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache MBs/sec	Cumulative number of host MBs read and written by the XtremSW Cache per second
XtremSW Cache Read Response Time	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write Response Time	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache Response Time	The average time it took the XtremSW Cache to serve one IO.
XtremSW Cache Avg Read Size (KB)	The average size of a read served by XtremSW Cache.

Table 38: Thin pool metrics (continued)

Metric	Description
XtremSW Cache Avg Write Size (KB)	The average size of a write served by XtremSW Cache.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	Number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: (Optimized MBs Read/sec * 1024 / Optimized Reads/sec)
Workload Score	A combined score metric that reflects the overall contribution of a volume or list of volumes to the array.
Front-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to front-end directors serving their IOs.
Back-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to back-end directors serving their IOs.
Disk Workload Score	A score metric that reflects the contribution of a volume or list of volumes to disks serving their IOs.

Volume metrics

This is a list of all KPIs that are available for volumes.

Table 39: Volume metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the Symmetrix volume.
Host Reads/sec	The total number of host read IO operations performed each second by the Symmetrix volume.
Host Writes/sec	The total number of host write IO operations performed each second by the Symmetrix volume.
Host Hits/sec	The total number of host read IO and write IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Read Hits/sec	The total number of host read IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Write Hits/sec	The total number of host write IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Misses/sec	The total number of host read IO and write IO operations missed each second by the Symmetrix volume.
Host Read Misses/sec	The total number of host read IO operations missed each second by the Symmetrix volume.
Host Write Misses/sec	The total number of host write IO operations missed each second by the Symmetrix volume.
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Host MBs Read/sec	A cumulative number of host MBs read per second.
Host MBs Written/sec	A cumulative number of host MBs written per second.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to the cache.
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to the cache.
Read Response Time (ms)	The average time it took the system to serve one read IO for this volume.
Write Response Time (ms)	The average time it took the system to serve one write IO for this volume.

Table 39: Volume metrics (continued)

Metric	Description
Read Miss Response Time (ms)	The average time a read miss operation was performed by the volume. A miss occurs when the requested data is not found in cache.
Write Miss Response Time (ms)	The average time a write miss operation was performed by the volume. A miss occurs when the write operation has to wait while data is destaged from cache to disks.
RDF/S Write Response Time (ms)	The average time it took the volume to serve one write IO.
% Reads	The percentage of IO operations that were reads.
% Writes	The percentage of IO operations that were writes.
% Read Hit	The percentage of read operations, performed by the Symmetrix volume, that were immediately satisfied by cache.
% Write Hit	The percentage of write operations, performed by the Symmetrix volume, that were immediately satisfied by cache.
% Read Miss	The percentage of read miss operations performed each second by the Symmetrix volume. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.
% Write Miss	The percent of write miss operations performed each second by the Symmetrix volume. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the Symmetrix volume that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the Symmetrix volume that were misses.

Table 39: Volume metrics (continued)

Metric	Description
Random IOs/sec	The number of IOs from a host not identified as part of a sequential stream.
Random Reads/sec	The number of read IO commands from a host not identified as part of a sequential stream.
Random Writes/sec	The number of write IO commands from a host not identified as part of a sequential stream.
Random Read Hits/sec	The number of random read IOs that were satisfied from the cache.
Random Write Hits/sec	The number of random write IOs that were immediately placed in cache because space was available.
Random Read Misses/sec	The number of random read IOs that were misses.
Random Write Misses/sec	The number of random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq ios per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$

Table 39: Volume metrics (continued)

Metric	Description
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total ios per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total ios per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total ios per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes read requests from the disk directors to cache.
BE Disk Read Response Time (ms)	The average time it takes read requests from the disk directors to cache.
BE Read Task Time (ms)	The time from the point when the HA puts the read request on the queue and the DA picks it up - can be considered queue time.
Num Invalid Tracks	Indicates internal copies in progress; this is an internal metric used by Engenuity.
Random Hits/sec	The total number of hits per second that were not sequential.
% Hit	The percent of operations that were immediately satisfied by cache.
% Miss	The percent of operations that were misses. A miss occurs when the operation cannot be immediately satisfied from cache because the data is not there or the operation has to wait while data is destaged from cache to disks.

Table 39: Volume metrics (continued)

Metric	Description
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
BE % Reads	The percent of read operations from the back-end directors that were immediately satisfied by cache.
BE % Writes	The percent of write operations from the back-end directors that were immediately satisfied by cache.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache per second.
BE Partial Sector Write Blocks	These metrics are for internal Symmetrix volume operations.
BE Optimize Writes (KB)	
BE XOR Reads	
BE XOR Reads (KB)	
BE Reads for Copy	
BE Writes for Copy	
BE Reads for Copy (KB)	
BE Writes for Copy (KB)	
BE Reads for VLUN Migration	
BE Reads for VLUN Migration (KB)	
BE Writes for VLUN Migration	
BE Writes for VLUN Migration (KB)	

Table 39: Volume metrics (continued)

Metric	Description
BE Writes for Rebuild	
BE Writes for Rebuild (KB)	
BE RDF Copy	
BE RDF Copy (KB)	
Device Capacity	The total capacity of the volume.
Device Block Size	The block size of the volume.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write.
XtremSW Cache IOs/sec	The average time it took XtremSW Cache to serve one IO.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
% XtremSW	The percent of XtremSW Cache IO that were reads.

Table 39: Volume metrics (continued)

Metric	Description
Cache Reads	
% XtremSW Cache Writes	The percent of XtremSW Cache IO that were writes.
% XtremSW Cache Read Hits	The percent of XtremSW Cache IO that were read hits.
XtremSW Cache MBs Read/sec	Cumulative number of host MBs read by the XtremSW Cache per second.
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache MBs/sec	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Read Response Time	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache Write Response Time	The average time it took the XtremSW Cache to serve one IO.
XtremSW Cache Response Time	The average size of a read served by XtremSW Cache.
XtremSW Cache Avg Read Size (KB)	The average size of a write served by XtremSW Cache.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	Number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: (Optimized MBs Read/sec * 1024 / Optimized Reads/sec)
Workload Score	A combined score metric that reflects the overall contribution of a volume or list of volumes to the array.

Table 39: Volume metrics (continued)

Metric	Description
Front-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to front-end directors serving their IOs.
Back-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to back-end directors serving their IOs.
Disk Workload Score	A score metric that reflects the contribution of a volume or list of volumes to disks serving their IOs.

Virtual pool tier metrics

This is a list of all KPIs that are available for Symmetrix virtual pool tiers.

Table 40: Virtual pool tier metrics

Metric	Description
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	The MBs transferred per second to the disk directors.
BE MBs Read/sec	The number of MBs read each second by the disk directors.
BE MBs Written/sec	The number of MBs written each second by the disk directors.
Read Miss Response Time (ms)	The average time that it took the Symmetrix system to serve one read miss IO for this pool.
Write Miss Response Time (ms)	The average time that it took the Symmetrix system to serve one write miss IO for this pool.
RDF/S Write Response Time (ms)	The average time it took the Symmetrix system to respond to the write request from the RDF/S group.
BE Read Request Time (ms)	The average time it took the back-end director to serve one read request.
BE Disk Read Response Time (ms)	The average time it takes cache to respond to a read request by the disk directors.
BE Read Task Time (ms)	The time from the point when the HA puts the read request on the queue and the DA picks it up - can be considered queue time.
% Used Capacity	The percent of the pools capacity that is used.
Total Pool Capacity (GB)	The total pool capacity in GBs.
Enabled Pool	The enabled pool capacity in GBs.

Table 40: Virtual pool tier metrics (continued)

Metric	Description
Capacity (GB)	
Used Pool Capacity (GB)	The used pool capacity in GBs.
Allocated Pool Capacity (GB)	The allocated pool capacity in GBs.
BE Read Response Time (ms)	A calculated value of the response time for all back-end read requests.
BE Write Response Time (ms)	A calculated value of the response time for all back-end write requests.
BE % Reads	Percentage of read operations from the back-end directors that were immediately satisfied by cache.
BE % Writes	Percentage of write operations from the back-end directors that were immediately satisfied by cache.
Ingress Tracks	The number of tracks entering the pool.
Egress Tracks	The number of tracks leaving the pool.
IO Density	The number of BE requests per GB of disk. (BE Reads + BE Writes) / allocated capacity With FAST moving active extents to higher tiers, this metric is a good indication of success (the IO density on Flash tiers should be higher than the density on SATA tiers.)
Total Tracks	The total number of tracks in the VP tier.
Compressed Tracks	The number of tracks in the VP tier that are compressed.
% Compressed Tracks	The percent of the total tracks that are compressed tracks.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read.
XtremSW Cache Writes /sec	The average time it took XtremSW Cache to serve one write.
XtremSW Cache IOs/sec	The average time it took XtremSW Cache to serve one IO.

Table 40: Virtual pool tier metrics (continued)

Metric	Description
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
% XtremSW Cache Reads	The percent of XtremSW Cache IO that were reads.
% XtremSW Cache Writes	The percent of XtremSW Cache IO that were writes.
% XtremSW Cache Read Hits	The percent of XtremSW Cache IO that were read hits.
XtremSW Cache MBs Read/sec	Cumulative number of host MBs read by the XtremSW Cache per second.
XtremSW Cache MBs Written/sec	Cumulative number of host MBs written by the XtremSW Cache per second.
XtremSW Cache MBs/sec	Cumulative number of host MBs read and written by the XtremSW Cache per second
XtremSW Cache Read Response Time	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write Response Time	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache Response Time	The average time it took the XtremSW Cache to serve one IO.
XtremSW Cache Avg Read Size (KB)	The average size of a read served by XtremSW Cache.
XtremSW Cache Avg Write Size (KB)	The average size of a write served by XtremSW Cache.

Table 40: Virtual pool tier metrics (continued)

Metric	Description
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Reads/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read/sec	Number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Size (KB)	Calculated value: (Optimized MBs Read/sec * 1024 / Optimized Reads/sec)
Workload Score	A combined score metric that reflects the overall contribution of a volume or list of volumes to the array.
Front-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to front-end directors serving their IOs.
Back-end Workload Score	A score metric that reflects the contribution of a volume or list of volumes to back-end directors serving their IOs.
Disk Workload Score	A score metric that reflects the contribution of a volume or list of volumes to disks serving their IOs.

Host IO Limits by SG metrics

This is a list of all KPIs that are available for Host IO limits set on front-end directors for the storage group.

Table 41: FE IO Limits by SG metrics

Metric	Description
Host IOs/sec	A host command for data transfer.
Read Response Time (ms)	The average time it takes the director/port to service one read IO for the SG.
Write Response Time (ms)	The average time it takes the director/port to service one write IO for the SG.
Host IO Limit Delayed IOs/sec	The sum of the IOs that were delayed due to the storage group's IO limits and bandwidth limits.
Host IO Limit Delayed Time (ms)	The time that IOs are delayed (due to the SG's set IO limits and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO limits and bandwidth limits).
Host IO Limit % Delayed IOs	The Host IO delayed IOs divided by the Host IO count.
Host IO Limit	The time that IOs are delayed (due to the SG's set IO limits and bandwidth

Table 41: FE IO Limits by SG metrics (continued)

Metric	Description
Average Delayed Time (ms)	limits) divided by the count of the IOs delayed (due to the SG's set IO limits and bandwidth limits).

Host IO Limits by FE metrics

This is a list of all KPIs that are available for front-end directors with Host IO limits.

Table 42: FE Limits by FE metrics

Metric	Description
Host IOs/sec	A host command for data transfer.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Host IO Limit IOs/sec	The number of IOs the FE port is performing on behalf of the specific SG.
Host IO Limit MBs/sec	The number of MBs the FE port is processing on behalf of the specific SG.
Host IO Limit Delayed IOs/sec	The sum of the IOs that were delayed due to the storage group's Host IO Limit.
Host IO Limit Delayed Time (ms)	The time that IOs are delayed (due to the SG's set IO and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO and bandwidth limits).
Host IO Limit % Delayed IOs	The quota delayed IOs divided by the IO count.
Host IO Limit Average Delayed Time (ms)	The time that IOs are delayed (due to the SG's set IO limits and bandwidth limits) divided by the count of the IOs delayed (due to the SG's set IO limits and bandwidth limits).

Settings

Exporting Performance Viewer settings

If you are using the offline Performance Viewer, you can export databases, thresholds, reports, user dashboards, and template dashboards.

To export Performance Viewer settings:

1. Select the Symmetrix system and **Performance**.
2. Select **Export PV Settings** from the **Common Tasks** panel.
You can export only the Symmetrix **SYMAPI database**, only the **Settings**, or both.
3. Enter a file name for the **SYMAPI database**.
4. Enter a file name for the **Settings**.
5. Click **OK**.

Unisphere attaches a date and time stamp to the name of the file(s) in GMT format, so that exported configuration files have *StartTime_FileName.zip* are automatically saved in the following directories:

SYMAPI Database: `SMAS\Backup\SPA\Viewer\SE`

Settings: `SMAS\Backup\SPA\Viewer\Settings`

Exporting and importing Performance settings

In large environments, it can be very time consuming to setup and configure multiple instances. Symmetrix systems may have very similar configuration and implementation's across the environment. To assist with the implementation process, the following settings can be exported and imported:

- ◆ The Performance metrics (KPIs) definitions with their default (not instance-specific) settings
- ◆ The user template dashboards

The settings are saved to an XML file. You select the the client side directory.

Before you begin:

For security compliance, a password must be supplied upon exporting and importing a performance settings file. If you pass the settings file to another administrator, make sure you also provide the password.

To export Performance settings:

1. Select a Symmetrix system.
2. Select any category (**Databases, Metrics, Reports, etc.**) under **Performance > Settings**.
3. Select **Export Settings** from the **Common Tasks** panel to open the Export Settings dialog box.

4. Select the settings to export: **Metrics Definitions, User Templates**, or both.
5. Enter and confirm a file password.

Note The settings file cannot be imported to another Unisphere environment without supplying the password set in this step.

6. Click **OK**. The **Save Settings** confirmation appears, click **OK**.
7. The Save File dialog box opens. Save the settings file to your preferred location.

To import Performance settings:

1. Select a Symmetrix system.
2. Select any category (**Databases, Metrics, Reports**, etc.) under **Performance > Settings**.
3. Select **Import Settings** from the **Common Tasks** panel to open the Import Settings dialog box.
4. Select the settings to import: **Metrics Definitions, User Templates**, or both.
5. Enter the file password (which was set during the export process). Click **OK**.
6. The Select File dialog box opens. Open the settings file from your preferred location.

File locations:

Unisphere places each file into its corresponding location on a new server setup:

- ◆ **catKpiMap.XML** under the `\Program`
Files\EMC\SMAS\jboss\standalone\data\spa\default-settings directory
The new metrics definition file will replace any existing metrics definition file.
- ◆ **SPA_TEMPLATE_DASHBOARD_0.XML** under the `\Program`
Files\EMC\SMAS\jboss\standalone\data\persistency\SPA_TEMPLATE_DASHBOARD directory
If a user template dashboards file already exists on the server, the existing templates will remain intact; only non-conflicting templates will be imported.

System Registrations

Registering Symmetrix systems

The Performance component of Unisphere for VMAX only collects performance data after you register your Symmetrix systems.

Only local Symmetrix systems can be registered for Performance data collection.

To register a Symmetrix system:

1. Select the Symmetrix system.
2. Click **Performance > Settings > Systems Registrations** to open the Systems Registrations view.
3. Click **Register** to open the Register dialog box. The Symmetrix ID automatically populates the **Select Symmetrix** list.
4. Select the type of performance data to collect, **Real Time, Diagnostic**, or both.
5. Click **OK**.

Changing data collection policies

To change collection policies:

1. Select **Performance** › **Settings** › **System Registrations** to open the System Registrations view.
2. Select the **Symmetrix ID** from the list of registered systems in the table.
3. **Click View Details.**
4. There are two data collection values that you can change:
 - **Collection Recovery (Hours)** — The amount of data that the storstopd daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
 - **Diagnostic Interval (Minutes)** — The frequency of polling in the Analyze Diagnostic view. The default 5 minutes is recommended.
5. Make your changes and click **Apply.**

Viewing data collection details

The performance data collection details show the default values for collection recovery and diagnostic interval. These default values can be adjusted.

To view data collection details:

1. Select **Performance** › **Settings** › **System Registrations** to open the System Registrations view.
2. Select the **Symmetrix ID** from the list of registered systems in the table.
3. **Click View Details.**

Viewing system registrations

To view system registrations:

1. Select **Performance** › **Settings** › **System Registrations** to open the System Registrations list view.

Use the **System Registrations** list view to view and manage Symmetrix system performance registration and data collection policies.

The following properties display:

- ◆ **Symmetrix ID** — The ID of the registered Symmetrix system.
- ◆ **Real Time** — An indicator that shows if the Symmetrix system is registered for Real Time data.
- ◆ **Diagnostic** — An indicator that shows if the Symmetrix system is registered for Diagnostic data.
- ◆ **Collection Recovery (Hours)** — The amount of data that the storstopd daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.

- ◆ **Diagnostic Interval (Minutes)** — The frequency of polling in the Analyze Diagnostic view. The default 5 minutes is recommended.
- ◆ **Server IP** — The IP address of the Unisphere for VMAX server.
- ◆ **SE Server IP** — The IP address of the Solutions Enabler server.

The following controls are available:

- ◆ **View Details** — See [Viewing system registration details](#) below.
- ◆ **Register** — See [System Registrations](#) on page 571.

Viewing system registration details

The system registration details view shows the default values for Symmetrix system registrations and data collection settings.

To view system registration details:

1. Select **Performance > Settings > System Registrations** to open the System Registrations view.
2. Select a **Symmetrix ID** and click **View Details** to open its **Details** view.

Properties panel

The following properties display:

- ◆ **Symmetrix ID** — The ID of the registered Symmetrix system.
- ◆ **Real Time** — An indicator that shows if the Symmetrix system is registered for Real Time data.
- ◆ **Diagnostic** — An indicator that shows if the Symmetrix system is registered for Diagnostic data.
- ◆ **Collection Recovery (Hours)** — The amount of data that the storstopd daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
- ◆ **Diagnostic Interval (Minutes)** — The frequency of polling in the Analyze Diagnostic view. The default 5 minutes is recommended.
- ◆ **Server IP** — The IP address of the Unisphere for VMAX server.
- ◆ **SE Server IP** — The IP address of the Solutions Enabler server.

The following controls are available:

- ◆ **Apply** — Saves any changes made to the properties.
- ◆ **Cancel** — Removes any changes made to the properties and returns the values to the defaults.

Metrics

Customize the metric list

All of the Symmetrix performance metrics are continuously collected and calculated. However, not all metrics are of equal importance to you. You can customize the list of metrics that display in the Performance dialogs and views.

To customize the metrics list:

1. Click **Performance > Settings > Metrics**. All metrics display.
2. To customize a metric, select the metric name and click **View Details**.
3. Select any of the following options for **Display Groups**:
 - a. Select **All** for any metric that you do not want to see by default, but may want to see occasionally.
 - b. Select **KPI** to display this metric by default. Selecting **KPI** automatically adds the metric to the **All** list.
 - c. Clearing both **KPI** and **All** boxes means the metric will not display for that category in Explore, Dashboards, Queries, and any other place metrics display.

All metrics always display in Metric Management.

Changing the threshold values


In addition, the Metric Manager allows for changing the default threshold values.

Threshold values are not required, but can be entered for any metric.

Default thresholds are predefined values assigned to a metric of a specific category. Changes to the default thresholds only affect the instances that are using the defaults.

Viewing metrics

Before you begin:

Optional: Select the Filter icon  and choose the **Category** you want to view.

To view metrics:

1. Select **Performance > Settings > Metrics** to open the **Metrics** list view.

Use the **Metrics** list view to view and manage metrics.

The following properties display:

- ◆ **Category** — The Symmetrix system component.
- ◆ **Metric** — The key performance indicator.
- ◆ **Display Group** — Indicates which metrics display for this component (**KPI**, **All**, or both).
- ◆ **First Threshold** — The default or custom first threshold.
- ◆ **Second Threshold** — The default or custom second threshold.

The following controls are available:

- ◆ **View Details** — Opens the **Details** view, from which you can [Metrics](#) (page 574).

Thresholds and Alerts

Creating a performance threshold and alert

You can use the default system values for thresholds and alerts, or create your own. When you set threshold values, you can optionally view them when you create charts for performance metrics in Analyze, Diagnostic view.

To create a threshold:

1. From the **Performance** section, select **Settings > Thresholds and Alerts** to open the Thresholds and Alerts list view.
2. Click **Create** to open the Create Performance Threshold and Alert dialog box.
3. Select the **Symmetrix ID, Category, Instance, and Metric**.
If the Category **BE, DX, or Disk** is selected, the **Instance** field defaults to **All**.
If there is a default value for the **Metric**, it automatically displays in the **Value** fields for **First Threshold** and **Second Threshold**.
4. Add a **Value** for the **First Threshold, Second Threshold, or both**.
5. *Threshold creation only:* Click **OK**.

To add an alert:

1. Click **Enable Alert** for **First Threshold, Second Threshold, or both**.
2. Select the **Severity** for the alert. Possible values are **Information, Warning, and Critical**.
3. Set the number of occurrences in the data samples which must happen before the alert is triggered. For example, if the threshold is breached 3 times out of 5 samples, initiate an alert.
4. Click **OK**.

Editing a performance threshold and alert

When you edit a threshold and alert setting, a symbol displays in the Custom column of the Alerts list to indicate the value has changed from the default.

A specific metric threshold can also be changed using the right-click menu from the table in the Analyze view.

To edit performance thresholds and alerts

1. From the **Performance** section, select **Settings > Thresholds and Alerts** to open the Thresholds and Alerts list view.
2. Select the item from the table and click **Edit** to open the Edit Performance Threshold and Alert dialog box.
3. Change the **Value** for **First Threshold, Second Threshold, or both**.
4. Select **Enable Alert** for **First Threshold, Second Threshold, or both**.

5. Set the number of occurrences in the data samples which must happen before the alert is triggered. For example, if the threshold is breached 3 times out of 5 samples, initiate an alert.
6. Click **OK**.

Deleting a performance threshold and alert

Before you begin:

Only custom values can be deleted. You cannot delete default thresholds.

To delete a performance threshold and alert:

1. From the **Performance** section, select **Settings > Thresholds and Alerts** to open the **Thresholds and Alerts** list view.
2. Select one or more rows and click **Delete**.
A confirmation message displays, "The threshold values will be reset to the default values and any alerts will be deleted."
3. Click **Delete**.

Viewing thresholds and alerts

1. Select the Symmetrix system.
2. Select **Performance > Settings > Thresholds and Alerts** to open the **Thresholds and Alerts** list.

The following properties display:

- ◆ **Symmetrix ID** — The Symmetrix system identifier.
- ◆ **Category** — The component category (**Array, Cache Partition, Disk**, etc.).
- ◆ **Instance** — Indicates whether the category or a specific instance of the category is measured.
- ◆ **Metric** — The key performance indicator (**IOs/sec, MBs/sec**, etc.).
- ◆ **First Threshold** — The user-defined first threshold.
- ◆ **Second Threshold** — The user-defined second threshold.
- ◆ **Custom** — Indicates (with a green circle) whether a custom threshold or alert has been set.

The following controls are available:

- ◆ **Create** — See [Thresholds and Alerts](#) on previous page.
- ◆ **Edit** — See [Editing a performance threshold and alert](#) on previous page.
- ◆ **Delete** — See [Deleting a performance threshold and alert](#) above.

Viewing thresholds on charts


The performance thresholds can be added to a chart in Analyze view by selecting **Display Thresholds** from the chart options.

Here is an example:

The Thresholds and Alerts list view shows the following thresholds for the % Hit metric on Symmetrix 0085:

Thresholds and Alerts						
Symmetrix ID	Category	Instance	Metric	First Threshold	Second Threshold	
000195700085	Array	000195700085	Host IOs/sec	⚠ 10000	❌ 2000000	
000195700085	Array	*	Host IOs/sec	1000000	2000000	
000195700085	Array	*	System WP Events/sec	1000	5000	
000195700085	Array	*	Host MBs/sec	0	0	
000195700085	Array	*	Avg Write Response Time (ms)	0	0	
000195700085	Array	*	% Hit	70	50	

In Analyze, Diagnostic view the Display Thresholds is turned on in the chart options. The red and green lines show the thresholds.



Real Time Traces

Creating a Real Time trace

A trace is a linear map; in Performance, a trace is one to four hours of real time data that is captured and saved. A trace can be configured in advance. Trace files can be retained up to 30 days.

Before you begin:

You must have access one or more Symmetrix systems with Real Time registrations.

To create a Real Time trace:

1. Select **Performance** > **Settings** > **Real Time Traces** to open the Real Time Traces view.
2. Click **Create** to open the Create Real Time Trace dialog box.
3. Select the **Symmetrix ID**. The Symmetrix IDs display if you are in **All Symmetrix** mode. If you have already selected a single Symmetrix system, this option is read-only.
4. Select the **Date** and **Start Time**. Click the calendar icon for the date selection and use the up and down arrows for setting the specific start time.
5. Select the **Number of Hourly Traces**. You can save up to 4 hours of data in four 1-hour files.

When configuring a new trace, the **Status** field displays the value **New**.

6. Select the **Retention** number of days. Values are 1–30.
7. *Optional:* Enter a **Description** of this trace.

All scheduled and executed traces display in the Real Time Traces view. The list can be for a specific Symmetrix system, or All Symmetrix systems.

The possible status values for the trace list are:

- ◆ **New** — Assigned while the new trace is being created.
- ◆ **Scheduled** — A trace has been configured and is scheduled to run at the specified start time.
- ◆ **Completed** — A scheduled trace has been completed successfully.
- ◆ **No Data** — No data for the scheduled trace was available at the time of execution.
- ◆ **Failed** — The execution of the trace failed. A system alert generates automatically.

Editing a Real Time Trace

A trace is 1 to 4 hours of real time data that is captured and saved. A trace can be configured in advance. Trace files can be retained up to thirty days.

Before you begin:

You must have access to licensed Symmetrix systems.

There must be a scheduled Real Time Trace showing in the Real Time Traces view. Only scheduled (not yet executed) traces can be modified.

To edit a trace:

1. Select **Performance** › **Settings** › **Real Time Traces** to open the Real Time Traces view.
2. Select a trace from the list and click **Edit** to open the Edit Real Time Trace dialog box.
3. Make changes to **Retention** and **Description**. These are the only values that can be changed.
4. Click **OK**.

Deleting a Real Time trace

Before you begin:

You must have one or more configured traces in the Real Time Traces list view.

To delete a Real Time trace:

1. Select **Performance** › **Settings** › **Real Time Traces** to open the **Real Time Traces** list view.
2. Select a trace from the list and click **Delete**.
3. Click **OK** on the confirmation message.

Viewing Real Time traces

A trace is a linear map; in Performance, a trace is one to four hours of real time data that is captured and saved.

A trace can be configured in advance. Trace files can be retained up to 30 days.

1. Select the Symmetrix system.
2. Select **Performance > Settings > Real Time Traces** to open the **Real Time Traces** list.

Use the **Real Time Traces** list view to display and manage Real Time traces.

The following properties display:

- ◆ **Symmetrix ID** — The Symmetrix system identifier.
- ◆ **Date** — The date of the scheduled trace.
- ◆ **Start Time** — The time of the scheduled trace.
- ◆ **Status** — The status of the scheduled trace. Possible values are:
 - **New** — The status assigned during creation.
 - **Scheduled** — The trace has been configured and scheduled to run at a specified time.
 - **Completed** — The scheduled trace has completed successfully.
 - **No Data** — No data was available for the scheduled trace at the time of execution.
 - **Failed** — Execution of the trace failed.
- ◆ **Retention (Days)** — The number of days the trace will be saved.
- ◆ **Description** — A user-defined description of the trace.

The following controls are available:

- ◆ **Create** — See [Real Time Traces](#) on page 577.
- ◆ **Edit** — See [Editing a Real Time Trace](#) on previous page.
- ◆ **Delete** — See [Deleting a Real Time trace](#) on previous page.

Reports

Creating performance reports

To create a performance report:

1. Select **Performance > Settings > Reports** to open the Reports list view.
2. Click **Create** to open the Create Report dialog box.
3. Enter a report **Name**.
4. Select the **Symmetrix ID**.
5. Select a **Format** (XML, CSV, BTP, UPV).

Note The UPV format is for [Creating reports for the Performance Viewer](#).

6. Select a **Type** (Diagnostic or Historical).
7. Select the **Time Range**. If you select **Custom**, a Time Selection dialog box displays for custom settings. The **Last Interval** option is only available for the **Diagnostic** type. This option executes a report on the array and director categories at each Diagnostic interval. In addition, every hour the system runs a cleanup job to remove Last Interval reports older than one hour.

Note The BTP format is not available for **Last Interval** reports.

8. *Optional:* Enter a report **Description**.
9. Make your **Schedule** selections.
10. Click **Add** to open the **Create Query Wizard**.

Note Reports must have at least one query.

11. Enter a query **Name**.
12. *Optional:* Enter a query **Description**.
13. Click **Next** to open page 2 of the wizard.
14. Select a **Category**, **Data Format**, and **Instances**. For example, **BE Director**, **Average**, and **Any instances in the selection**.
15. Click **Next** to open page 3 of the wizard.
16. Select the **Metrics** (for the category/data format/instance you selected on page 2 of the wizard). Use multi-select for multiple metrics.
17. Click **Finish** to close the wizard.
18. Click **OK**.

Copying performance reports

Before you begin:

There must be an existing report in the Reports list view.

To copy a performance report:

1. Select **Performance > Settings > Reports** to open the Reports list view.
2. Select a report and click **Copy** to open the Copy Report dialog box.
3. Make your changes to the report **Name**, **Symmetrix ID**, **Format**, **Type**, **Time Range**, and **Description**.
4. *Optional:* Select **Recurring** to automatically schedule this report.

Note You cannot select **Recurring** for custom time selections.

5. Select one of the following actions in the **Queries** section:
 - a. Click **Add** to open the Create Query Wizard and create a new query.
 - b. Select a query and click **Delete**. Click **OK** on the confirmation message.
 - c. Select a query and click **Edit** to open the Edit Query wizard.

Note Reports must have at least one query.

6. Make your changes in each page of the wizard.
7. Click **Finish** to close the wizard.
8. Click **OK**.

For more information about reports, refer to [Reports](#) on previous page.

Editing performance reports

Before you begin:

There must be an existing report in the Reports list view.

To edit a performance report:

1. Select **Performance** > **Settings** > **Reports** to open the Reports list view.
2. Select a report and click **Edit** to open the Edit Report dialog box.
3. Make your changes to the report **Symmetrix ID, Format, Type, Time Range, and Description**.

Note You cannot change the report name.

4. *Optional:* Schedule the report.
5. Select one of the following actions in the Queries section:
 - a. Click **Add** to open the Create Query Wizard and create a new query.
 - b. Select a query and click **Delete**. Click **OK** on the confirmation message.
 - c. Select a query and click **Edit** to open the Edit Query wizard.

Note Reports must have at least one query.

6. Make your changes in each page of the wizard.
7. Click **Finish** to close the wizard.
8. Click **OK**.

For more information about reports, refer to [Reports](#) on page 579.

Scheduling performance reports

Before you begin:

You must have a configured performance report in the **Reports** list.

To run a performance report:

1. Select **Performance** > **Settings** > **Reports** to open the Reports list.
2. Select one or more reports and click **Schedule** to open the Schedule Report dialog box.
3. Select a **Execution Time** and Days to Run.
4. Click **OK**.

Note Scheduled reports save to this default location:
`SMAS\jboss\standalone\data\spa\querydata\`

Running performance reports

Before you begin:

You must have a configured performance report in the **Reports** list.

To run a performance report:

1. Select **Performance** › **Settings** › **Reports** to open the **Reports** list.
2. Select one or more reports and click **Run Now**.
3. Click **OK** on the **Save Query Results** confirmation.
4. Enter a location in the **Select location for download** dialog box, enter a file name and extension (if necessary), and click **Save**.

The report saves to the selected location.

Canceling a scheduled report

Before you begin:

You must have a scheduled performance report in the Reports list view.

To cancel a scheduled report:

1. Select **Performance** › **Settings** › **Reports** to open the Reports list view.
2. Select one or more reports and click **Cancel Schedule**.
3. Click **OK** in the confirmation message.

Deleting performance reports

Before you begin:

There must be an existing report in the Reports list view.

To view performance reports:

1. Select **Performance** › **Settings** › **Reports** to open the Reports list view.
2. Click **Delete**.
3. Click **OK** on the confirmation message.

Viewing reports

1. Select the Symmetrix system.
2. Select **Performance** › **Settings** › **Reports** to open the **Reports** list.

Use the Reports list to display and manage reports.

The following properties display:

- ◆ **Name** — The user-defined report name.
- ◆ **Symmetrix ID** — The Symmetrix system identifier.
- ◆ **Format** — The report format (XML, CSV, BTP, UPV).
- ◆ **Recurring** — Indicates whether the report will run automatically.
- ◆ **Run Date** — The date of the report.
- ◆ **Description** — The user-defined description.

The following controls are available:

- ◆ **Create** — See [Reports](#) on page 579.
- ◆ **Copy** — See [Copying performance reports](#) on page 580.
- ◆ **Edit** — See [Editing performance reports](#) on page 581.
- ◆ **Schedule** — See [Scheduling performance reports](#) on page 581.
- ◆ **Cancel Schedule** — See [Cancelling a scheduled report](#) on previous page.
- ◆ **Run Now** — See [Running performance reports](#) on page 581.
- ◆ **Delete** — See [Deleting performance reports](#) on previous page.

Note Scheduled reports save to this default location:
`SMAS\jboss\standalone\data\spa\querydata\`

Viewing configured performance reports

This section describes the performance reports view.

To view performance reports:

1. Select **Performance > Settings > Reports** to open the **Reports** list.

From the Reports list view you can:

[Reports](#) (page 579)

[Running performance reports](#) (page 581)

[Scheduling performance reports](#) (page 581)

[Editing performance reports](#) (page 581)

[Copying performance reports](#) (page 580)

See [Cancelling a scheduled report](#) on previous page.

See [Deleting performance reports](#) on previous page.

When you select **Run Now**, you choose the location for the generated reports.

Scheduled reports automatically save to this default location:

`SMAS\jboss\standalone\data\spa\querydata\`

Databases

Backing up a database

The backup database option is available for one or more Symmetrix systems, regardless of their registration status. By default, only Historical data is backed up.

To back up the performance database:

1. Select **Performance** › **Settings** › **Databases** to open the Databases view.
2. Select a database from the list and click **Backup** to open the Database Backup wizard.
3. Select the **Symmetrix** system.
4. Enter a backup **File Name**.
5. *Optional:* Select any of the following **Data to Save**:
 - **Last day of Diagnostics**
 - **Named Real Time Traces**

Note **Historical** is always selected and cannot be unselected.

6. Set the **Execution Time**. Possible values are **Now**, and **Schedule** (specific date and time).
 - **Now** — Go to step 7.
 - **Schedule only**— Click **Next**. Select the **Start Date** and time. Select if you want this to be a **Recurring** database backup. If **Yes**, select the days on which the database backups should occur.
7. Click **Finish**.

A database backup can take some time.

Also, refer to See [Removing database backup files](#) on the facing page..

Canceling a scheduled database backup

Before you begin:

There must be one or more database backups scheduled for a future time.

To cancel a scheduled database backup:

1. Select **Performance** › **Settings** › **Databases** to open the Databases view.
2. Select a database and click **Cancel Schedule**.
3. Click **OK** on the confirmation message.

Restoring a database

You can restore a Symmetrix system's performance database regardless of whether it is registered.

Before you begin:

There must be an existing database file.

To restore a performance database:

1. Select **Performance, Settings, Databases** to open the **Databases** list.
2. Select a database from the list and click **Restore** to open the Restore Database dialog box.
3. Select the **Symmetrix** system.
4. Select the **Backup File**.
5. *If available*: Select if you want to restore Diagnostic data and if you want to restore Real Time traces.
6. Click **OK**.

Deleting a database

The database delete action can only be performed on Symmetrix systems that are not registered. The delete action removes all references to the Symmetrix systems in the master database and removes the data. Backup files are not removed.

Before you begin:

You must have one or more Symmetrix systems that are no longer registered and have a performance database.

To delete a performance database:

1. Select **Performance > Settings > Databases** to open the Databases view.
2. Select a database from the list and click **Delete**.
3. Select **Yes** or **No** to delete backup files associated with the database.
4. Click **OK**.

Removing database backup files

This procedure is for removing unneeded database backup files.

To delete database backup files:

1. Select a Symmetrix system
2. Select **Performance > Settings > Databases** to open the **Databases** list view.
3. Select a database and click **Remove Backup Files** to open the **Remove Backup Files** dialog box.
4. Select the file(s) to delete from the list of backup files and click **OK**.

Viewing databases

To view databases:

1. Select **Performance > Settings > Databases** to open the **Databases** list.

Use the **Databases** list view to view and manage databases.

The following properties display:

- ◆ **Symmetrix ID** — The Symmetrix system identifier.
- ◆ **Status** — The status of the performance database.
- ◆ **Task Status** — The status of a database action (backup, restore).
- ◆ **Oldest Historical Data** — Date and time of the earliest Historical data.
- ◆ **Oldest Diagnostics Data** — Date and time of the earliest Diagnostic data.
- ◆ **Last Backup** — The last backup taken (in Days).
- ◆ **Last Retention** — The last time data was saved (in hours).
- ◆ **Scheduled Backup** — Time of scheduled backup.

The following controls are available:

- ◆ **Backup** — See [Databases](#) on page 584.
- ◆ **Restore** — See [Restoring a database](#) on page 584.
- ◆ **View Details** — See [Viewing database details](#) below.
- ◆ **Cancel Schedule** — See [Canceling a scheduled database backup](#) on page 584.
- ◆ **Delete Database** — [Deleting a database](#)

Viewing database details

Before you begin:

There must be one or more Symmetrix system performance databases.

To view database details:

1. Select **Performance, Settings, Databases** to open the **Databases** list.
2. Select a database and click **View Details**.

The following details display:

- ◆ **Symmetrix ID** — The Symmetrix system identifier.
- ◆ **Registered to Collection** — Shows the performance views available. Possible values are **Real Time** and **Diagnostic**, or **Not Registered**.
- ◆ **Database Administration**
 - **Last backup time** — Date and time of last backup.
 - **Last restore time** — Date and time of last restore.
 - **Last aggregation time** — Last date and time that historical data was derived from diagnostics data.
 - **Last retention time** — Date and time of last retention.
 - **Last real time retention time** — Date and time of last Real Time retention.
- ◆ **Historical**

- **Oldest available data** — Date and time of first available Historical data.
- **Latest available data** — Date and time of latest Historical data.
- **Set Historical data retention (months)** — Number of months to save Historical data. Possible values are 6 to 36 months. This option can be changed. After any change, click **Apply**.
- ◆ **Diagnostics**
 - **Oldest available data** — Date and time of first available diagnostic data.
 - **Latest available data** — Date and time of latest diagnostic data.
 - **Set Diagnostics data retention (days)** — Number of days to save Diagnostic data. Possible values are 1 to 7 days. This option can be changed. After any change, click **Apply**.
- ◆ **Backup Files**
 - **Set Backup file retention (number of files)** — Number of backup files to save. When this number is exceeded, the oldest file is automatically deleted. This option can be changed. After any change, click **Apply**.
 - **Number of existing backup files** — Number of existing backup files for this database.
- ◆ **Scheduled Backup Execution Options**
 - **Next Backup** — Displays the date/time for the next scheduled backup (if any).
 - **Recurring** — Displays the days (Monday, Tuesday, etc) set for a recurring backup.
 - **Last Day of Diagnostics** — Indicates whether the last day of diagnostic statistics will be included in the next database backup.
 - **Named Real Time Traces** — Indicates whether any named real time traces will be included in the next database backup.

The following controls are available:

Apply — Saves any changes made to database details.

Cancel — Clears any changes made to database details.

