EMC® Unisphere for VMAX

VERSION 8.1

Online Help

REV 01

EMC²

Copyright © 2011 - 2015 EMC Corporation. All rights reserved. Published in the USA.

Published September, 2015

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.

EMC^2, 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: Introduction to Unisphere for VMAX **19**

Chapter 2: Getting Started **21**

Operating as the initial setup user 22
Touring the interface 23
Using the Home Dashboard 27
Using the system selector 29
Discovering storage systems 30
Exporting data 30
Refreshing console information 30
Refreshing storage system information 31
Setting preferences 31
Searching for storage objects 32
Filtering lists of objects 33
Exiting the console 35
Getting help 36

Chapter 3: Administration **37**

Setting system preferences 38
Displaying login messages 41
Database backup 42
Backing up the database server 42
Viewing database backups 43
Deleting database backups 43
Alert settings 44
Alerts 44
Alert policies 72
Pool threshold alerts 74
Performance thresholds and alerts 76
SLO compliance alert policies 80
Runtime checks 83
Security 85
Understanding user authorization 85
Viewing user sessions 87

Authentication 87
Authorization 90
Local Users 93
Link and launch 101
Creating link-and-launch client registrations 101
Editing link-and-launch client registrations 101
Deleting link-and-launch client registrations 102
Viewing link and launch client registrations 102

Chapter 4: Storage Management 103
VMAX provisioning 104
Provisioning VMAX storage (HYPERMAX OS 5977 and higher) 104
Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) 106
Provisioning VMAX storage (Enginuity 5874 - 5876) 108
Using the Provision Storage wizard (Enginuity 5874 - 5876) 110
Reference workloads 114
Storage groups 117
Managing storage groups (HYPERMAX OS 5977 or higher) 117
Creating storage groups (HYPERMAX OS 5977 or higher) 119
Creating storage groups (Enginuity 5874 - 5876) 120
Adding volumes to storage groups 121
Copying volumes between storage groups 122
Moving volumes between storage groups 123
Removing volumes from storage groups 124
Expanding storage groups (Enginuity 5874 - 5876) 126
Modifying storage groups 129
Moving storage groups to EMC CloudArray 131
Protecting storage groups 132
Converting storage groups to cascaded 140
Changing Storage Resource Pools for storage groups 140
Associating FAST policies and storage groups 141
Disassociating FAST policies and storage groups 143
Reassociating FAST polices and storage groups 144
Adding or removing cascaded storage groups 145
Deleting storage groups 145
Setting host I/O limits 146
Managing VP compression on storage groups 148

Viewing storage groups 148
Fully Automated Storage Tiering 165
Understanding FAST (HYPERMAX OS 5977) 165
Understanding FAST (Enginuity 5874 - 5876) 176
Understanding Workload Planner 199
Volume management 208
Managing volumes (HYPERMAX OS 5977 or higher) 208
Managing volumes (Enginuity 5671 - 5876) 208
Creating volumes 210
Deleting volumes 221
Duplicating volumes 221
Assigning Symmetrix priority 223
Changing volume configuration 223
Expanding existing volumes 225
Mapping volumes 226
Unmapping volumes 227
Setting optimized read miss 227
Setting volume status 229
Setting volume attributes 230
Setting volume identifiers 231
Setting volume names 231
Setting replication QoS 232
Managing Meta Volumes 235
Disk groups 245
Renaming disk groups 245
Removing disks from disk groups 245
Deleting disk groups 245
Viewing disk groups 246
Viewing disk group details 247
Viewing disks in disk group 248
Viewing disk details 248
Virtual Provisioning 251
DATA volumes 251
Thin pools 257
Thin volumes 271
Understanding Virtual LUN Migration 277
Migrating regular storage group volumes 277

Migrating regular volumes 278
Migrating thin storage group volumes 279
Migrating thin volumes 279
Terminating a VLUN migration session 280
Viewing VLUN migration sessions 281
External Storage 282
Viewing external storage 283
Rescanning external storage 284
Understanding ProtectPoint 284
Understanding FAST.X 288
Filtered results table 289
Understanding Federated Tiered Storage 293
Understanding storage templates 296
Creating storage templates 296
Exporting storage templates 298
Importing storage templates 298
Deleting storage templates 298
Viewing storage templates 299
Viewing storage template details 300
Reservations 302
Reserving volumes 302
Adding volumes to reservations 303
Removing volumes from reservations 303
Releasing reservations 304
Viewing reservations 304
Viewing reservation details 304
Understanding Optimizer 306
Managing Optimizer 306
Starting/Stopping Optimizer 307
Enabling/Disabling Optimizer 308
Locking/Unlocking Optimizer 308
Approving Optimizer Swaps 309
Viewing Optimizer swap/move lists 310
Viewing Optimizer swap/move history 310
Managing vVols 312
Viewing storage containers 313
Viewing storage resource details 313

Creating storage containers 314
Deleting storage containers 315
Viewing storage resources 315
Viewing storage resource details 316
Viewing storage resource related SRPs 316
Adding storage resources to storage containers 317
Removing storage resources from storage containers 317
Viewing protocol endpoints 317
Viewing protocol endpoint details 318
Provisioning protocol endpoints to hosts 318
Viewing storage container alerts 319

Chapter 5: Host Management 321

Hosts 322
Creating hosts 322
Adding initiators to hosts 323
Removing initiators from hosts 324
Modifying hosts 324
Renaming hosts/host groups 325
Setting host/host group flags 326
Deleting hosts/host groups 326
Viewing hosts/host groups 326
Viewing host/host group details 327
Viewing host initiators 329
Host/Host group flags 329
Host Groups 331
Creating host groups 331
Adding hosts to host groups 332
Removing hosts from host groups 332
Modifying host groups 333
Renaming hosts/host groups 334
Setting host/host group flags 334
Deleting hosts/host groups 334
Viewing hosts/host groups 335
Viewing host/host group details 336
Viewing hosts in host groups 337
Masking views 338

Creating masking views 338
Renaming masking views 338
Deleting masking views 339
Viewing masking views 339
Viewing masking view connections 339
Viewing masking view details 341
Initiators 343
Masking volumes 343
Unmasking volumes 344
Setting initiator port flags 344
Setting initiator attributes 344
Renaming initiator aliases 345
Replacing initiators 346
Removing masking entries 346
Viewing initiators 346
Viewing initiator details 347
Port groups 350
Creating port groups 350
Deleting port groups 351
Adding ports to port groups 351
Removing ports from port groups 352
Renaming port groups 352
Viewing port groups 353
Viewing port groups details 353
Viewing ports in port group 355
Viewing port details 356
Virtual servers 357
Adding a new virtual server 357
Adding storage to a VM 357
Removing a virtual server 358
Removing storage from a VM 358
Changing the password on a virtual server 359
Viewing virtual servers 360
Viewing the details of a virtual server 360
CU images 362
Mapping CKD volumes 362
Unmapping CKD volumes 362

Assigning alias addresses 363
Removing alias addresses 364
Assigning alias ranges 364
Removing alias ranges 365
Assigning alias counts 365
Removing alias counts 366
Viewing CU images 366
Viewing CU image details 367
Host aliases 369
Masking volumes 369
Creating host aliases 370
Adding initiators to host aliases 370
Removing initiators from host aliases 370
Unmasking volumes from host aliases 370
Deleting host aliases 371
Renaming host aliases 371
Viewing host aliases 371
Viewing host alias details 372
Host Cache Adapters 373
Viewing host cache adapters 373

Chapter 6: Data Protection 375
Monitoring protection 376
Protection Dashboard: 376
Local Replication 379
TimeFinder/Clone 379
TimeFinder/Snap 395
TimeFinder SnapVX 408
TimeFinder/Mirror 423
Device Groups 433
Creating device groups 433
Adding volumes to device groups 434
Removing volumes from device groups 434
Enabling consistency protection 434
Disabling consistency protection 435
Renaming device groups 436
Deleting device groups 436

Viewing device groups 437
Viewing device group details 437
Remote Replication 441
Managing remote replication sessions 441
SRDF 441
Properties panel 450
Related Object panel 453
Related Object panel 455
Related Object panel 457
Understanding Virtual LUN Migration 484
Migrating regular storage group volumes 484
Migrating regular volumes 485
Migrating thin storage group volumes 486
Migrating thin volumes 486
Viewing VLUN migration sessions 487
Viewing VLUN migration session details 487
Replication Groups and Pools 489
SRDF/A DSE Pools 489
TimeFinder Snap Pools 493
SRDF Groups 496
Understanding RecoverPoint 517
Tagging and untagging volumes for RecoverPoint 518
Viewing RecoverPoint sessions 519
Viewing RecoverPoint session details 519
Viewing RecoverPoint tagged volumes 520
Viewing RecoverPoint tagged volume details 520
Protecting storage groups using RecoverPoint 522
Viewing RecoverPoint volumes 524
Viewing RecoverPoint clusters 526
Viewing RecoverPoint cluster details 526
Viewing RecoverPoint splitters 527
Viewing RecoverPoint appliances 528
RecoverPoint systems 529
RecoverPoint consistency groups 532
RecoverPoint copies 534
RecoverPoint replication sets 537
RecoverPoint links 538

Open Replicator 541
Creating Open Replicator copy sessions 541
Managing Open Replicator sessions 542
Activating Open Replicator sessions 543
Restoring Open Replicator sessions 543
Renaming Open Replicator sessions 544
Removing Open Replicator sessions 544
Setting Open Replicator session background copy mode 544
Setting Open Replicator session donor update off 545
Setting Open Replicator session front end zero detection off 545
Setting Open Replicator session pace 545
Setting Open Replicator ceiling 546
Terminating Open Replicator sessions 546
Viewing Open Replicator sessions 546
Viewing Open Replicator session details 547
Open Replicator session options 548
Open Replicator flags 550
Understanding Federated Live Migration 552
Setting up/Running Federated Live Migration 553
Creating a FLM session 555

Chapter 7: Performance management 559
Common Tasks 560
Exporting Performance settings 560
Importing Performance settings 560
Exporting Performance Viewer settings 561
About exporting and importing Performance settings 561
Saving dashboards and charts 562
Filtering performance data 563
Changing the time range 563
Monitor view 566
Managing dashboards 566
Viewing dashboards 568
Creating a dashboard or template folder 569
Creating a dashboard with charts 569
Copying a dashboard 569
Editing a template dashboard 570

Editing a dashboard 570
Deleting a dashboard 570
Running a report from the dashboard 571
Saving a dashboard as a template 571
Saving dashboard changes 571
Managing EMC Views 572
Navigating from a heatmap to an analyze view 572
Scheduling a report from the dashboard 572
Using EMC summary dashboards 573
Analyze view 578
Navigating in the Analyze tab 578
Viewing information in the Analyze view 579
Creating a dashboard from an Analyze view 579
Creating a template dashboard from an Analyze view 579
Viewing a Real Time trace 580
Real Time Analysis views 580
Root Cause Analysis views 583
Trending & Planning views 597
Charts view 613
Using the UI elements in the Charts view 613
Creating charts 615
Customizing charts 616
Customizing the Charts view 618
Plan view 620
Using projection dashboards 620
Reports view 623
Viewing performance reports 623
Creating performance reports 623
Creating queries using the Create Query wizard 624
Copying performance reports 625
Editing performance reports 626
Scheduling performance reports 627
Cancelling a scheduled report 627
Running performance reports 627
Deleting performance reports 628
Settings 629
System Registrations 629

Performance Thresholds and Alerts 631
Real Time Traces 634
Databases 636
Metrics 640
Metrics 642
Array metrics 642
BE Director (DA) metrics 647
BE Emulation metrics 649
BE Port metrics 650
Board metrics 651
Cache Partition metrics 652
Core metrics 655
DATA Volume metrics 656
Database metrics 658
Database by Pool metrics 664
Device Group metrics 665
Disk metrics 672
Disk Bucket metrics 674
Disk Group metrics 675
Disk Group tier metrics 677
Disk Technology metrics 679
DSE Pool metrics 681
DX Emulation metrics 682
DX Port metrics 683
EDS Director metrics 684
EDS Emulation metrics 685
External Director metrics 686
External Disk metrics 689
External Disk Group metrics 690
FAST VP Policy metrics 691
FE Director metrics 692
FE Emulation metrics 695
FE Port metrics 696
FE Port - FE metrics 697
FE Port - SG metrics 698
FICON Emulation metrics 699
Host metrics 700

Host IO Limit by FE metrics 701
Host IO Limit by SG metrics 702
IM Director metrics 703
IM Emulation metrics 704
Initiator metrics 705
Initiators by Port metrics 706
IP Interface metrics 707
iSCSI Target metrics 708
Masking View metrics 709
Metas metrics 710
Other - Pool Bound Volume metrics 718
Pool by Storage Group metrics 726
Port Group metrics 727
RDF Director metrics 728
RDF Emulation metrics 730
RDF Port metrics 731
RDF/A Group metrics 732
RDF/S Group metrics 737
SAVE Volume metrics 742
Snap Pool metrics 750
Spare Disk metrics 751
SRP metrics 753
Storage Group metrics 754
Storage Group by Pool metrics 762
Storage Group by Tier metrics 763
Thin Pool metrics 764
Thin Volume metrics 767
Tier by Storage Group metrics 773
Virtual Pool Tier metrics 774
Volume metrics 776

Chapter 8: Databases 785

Understanding Database Storage Analyzer 786
Viewing databases 787
Viewing database details 788
Launching Database Storage Analyzer 789

Chapter 9: System Management **791**
Monitoring storage systems 792
System Dashboard: 792
Viewing storage system details 795
Setting system attributes 799
Setting CPU I/O resource distribution 802
Understanding eNAS 803
Discovering eNAS control stations 803
Managing File Storage 803
Provisioning storage for file 805
Launching Unisphere for VNX 808
Managing file storage groups 809
Managing file masking views 810
Viewing file systems 811
Viewing file system details 811
Viewing file system storage pools 812
Viewing file system storage pool details 813
Managing file storage alerts 815
Audit log 819
Viewing the system audit log 819
Viewing Symmetrix audit log details 819
Filtering audit log records 820
Hardware management 822
Converting directors 822
Associating directors with ports 823
Setting director port attributes 823
Associating directors and ports 825
Disassociating directors and ports 826
Enabling and disabling director ports 826
Performing system health checks 826
Naming storage systems 827
Locating VMAX systems 827
Replacing failed drives 828
Viewing hardware components 831
Managing jobs 848
Making configuration changes safely 849

Understanding task persistence 849
Previewing 850
Scheduling jobs 851
Running jobs 851
Rescheduling jobs 851
Renaming jobs 851
Reordering tasks within a job 852
Grouping jobs 852
Un-grouping jobs 852
Stopping jobs 852
Deleting jobs 853
Viewing the job list 853
Viewing job details 855
Understanding licenses 857
Licenses 857
Unisphere for VMAX licensing 857
Installing licenses 857
Removing host-based licenses 858
Viewing Symmetrix entitlements 858
Viewing host-based licenses 860
Viewing license usage 860
Understanding access controls 862
Opening access controls 862
Creating access groups 862
Adding access ID to access groups 863
Removing access IDs from access groups 864
Deleting access groups 864
Viewing access groups 864
Creating access pools 864
Modifying access pools 865
Deleting access pools 865
Viewing access pools 866
Modifying access types 866
Creating access control entries 866
Deleting access control entries 867
Viewing access control entries 867
Viewing access control entry details 867

Viewing access groups 868
Viewing access group details 869
Viewing access IDs 869
Viewing access pools 870
Viewing access pool details 870
Viewing access pool volumes 871
Viewing access types 871
Access types 872
Understanding dynamic cache partitioning 874
Enabling/Disabling dynamic cache partitioning 874
Creating dynamic cache partitions 874
Assigning dynamic cache partitions 875
Deleting dynamic cache partitions 875
Running in analyze mode 876
Viewing dynamic cache partitions 876
Viewing dynamic cache partition details 877
iSCSI 879
Viewing the iSCSI dashboard 880
Creating an iSCSI target 883
Creating an IP interface 885
Adding an IP route 886
Deleting an iSCSI target 886
Deleting an IP interface 887
Removing an IP Route 887
Attaching an IP interface to an iSCSI target 888
Attaching an iSCSI target to an IP interface 890
Detaching an IP interface from an iSCSI target 891
Disabling an iSCSI target 891
Enabling an iSCSI target 891
Setting port flags 892
Viewing the iSCSI directors list 893
Viewing the iSCSI director details 894
Viewing the IP interfaces list 895
Viewing the IP interfaces details 896
Viewing the iSCSI targets list 899
Viewing the iSCSI target details 900
Viewing the IP routes list 901

Viewing the IP routes details 902
Viewing the iSCSI ports list 903
Viewing the iSCSI ports details 904

Introduction to Unisphere for VMAX

Unisphere for VMAX is a web-based application that allows you to configure and manage VMAX storage systems. You can use Unisphere to:

- Manage user accounts and roles
- Perform configuration operations (create thin volumes, mask volumes, set storage attributes, set volume attributes, and set port flags)
- Manage volumes (change volume configuration, set volume status, and create/dissolve meta volumes)
- Perform and monitor replication and backup operations:
  - TimeFinder® SnapVX
  - TimeFinder/Snap
  - TimeFinder VP Snap
  - TimeFinder/Clone
  - TimeFinder/Mirror
  - Symmetrix Remote Data Facility (SRDF®)
  - Open Replicator for Symmetrix (ORS)
  - ProtectPointTM
- Manage advanced storage features, such as:
  - Fully Automated Storage Tiering (FAST)TM
  - Fully Automated Storage Tiering for virtual pools (FAST VP)
  - Service Level Objective provisioning
  - Workload planning
  - FAST Array Advisor
  - Enhanced Virtual LUN Technology
  - Auto-provisioning Groups
  - Virtual Provisioning
  - Federated Live Migration (FLM)
  - Federated Tiered Storage (FTS)
  - Embedded NAS (eNAS)
  - FAST.X
- Monitor alerts, including the ability to configure external alert notifications
- Monitor storage system performance data:
  - Monitor performance and capacity over time

- Drill-down through data to investigate issues
- View graphs detailing system performance
- Set performance thresholds and alerts
- View high frequency metrics in real time
- Perform root cause analysis
- View storage system heatmaps
- Execute scheduled and ongoing reports (queries), and export that data to a file
- Utilize predefined dashboards for many of the system components
- Customize your own dashboard templates
- Execute scheduled export of performance dashboards

◆ Monitor and troubleshoot database performance issues using Database Storage Analyzer

CHAPTER 2
Getting Started

This chapter contains the following:

Operating as the initial setup user 22
Touring the interface 23
Using the Home Dashboard 27
Using the system selector 29
Discovering storage systems 30
Exporting data 30
Refreshing console information 30
Refreshing storage system information 31
Setting preferences 31
Searching for storage objects 32
Filtering lists of objects 33
Exiting the console 35
Getting help 36

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 storage systems that do not have defined roles (authorization rules). Once an Administrator or Security is assigned to a storage 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:

- Local Users on page 93
- Authorization on page 90

Touring the interface

The Unisphere for VMAX interface consists primarily of dashboards and views. Dashboards provide a high-level view of your storage environment, and views provide a low-level view of the various storage objects (groups, volumes, etc.). While dashboards provide some management functionality, most of your management tasks will start in the views.

The following example shows a typical view. In this case, the **Volumes** list view for a storage group.



Title bar:

Clicking on this displays a drop down menu. The menu options available are: Help, About, Export, Settings and Logout. The drop down menu also displays information about the user who is logged onto Unisphere.

Opens the search dialog.

Menu bar:

Note: Depending on your storage operating environment, some of the following options may not apply.

Home The **Home** menu provides access to the following operations:

- View the **Home Dashboard**, from which you can view how the well the storage systems are performing, including capacity usage, response time, SLO compliance, and alerts.
- View and manage Administration settings, including:
  - Alert policies, thresholds, and notifications
  - Authentication options (local directory, Windows OS/AD, LDAP-SSL)
  - System preferences
  - User authorizations
  - Link and launch client registrations
- View, acknowledge, and delete alerts for all visible storage systems
- View/Export storage system capacity information

System selector The system selector displays a list of all managed storage systems. It allows you to toggle the interface between two modes, All Symmetrix and individual systems. Selections you make in the list provide the context for subsequent operations. For

more information, refer to Using the system selector (page 29).

Storage

The **Storage** menu provides access to the following operations, depending on the storage environment:

- View and manage storage groups (Enginuity version 5874 - 5876)
- Provision storage to hosts
- View and mange service levels (HYPERMAX OS 5977 or higher)
- View and manage Storage Resource Pools (HYPERMAX OS 5977 or higher)
- View and manage volumes
- View and manage external storage (Enginuity 5876, HYPERMAX OS 5977 or higher)
- View and manage the FAST™ controller policies (Enginuity 5874 - 5876)
- View and manage thin pools (Enginuity 5773 - 5876)
- View and manage storage tiers (Enginuity 5874 - 5876)
- View and manage storage volumes
- View and manage storage templates (Enginuity version 5874 - 5876)
- View and manage disk groups
- View and manage Symmetrix Optimizer (Enginuity version 5671 - 5876)

The **Storage** menu is only available when you select an individual storage system from the system selector.

Hosts

The **Hosts** menu provides access to the following operations:

- View and manage hosts
- View and manage masking views
- View and manage initiators (rename alias, set attributes, set flags, replace initiator)
- View and manage port groups
- View XtremSW cache adapters
- View and manage virtual servers

The **Hosts** menu is only available when you select an individual storage system from the system selector.

Data Protection

The **Data Protection** menu provides access to the following operations:

- View and manage TimeFinder sessions
- View and manage SRDF sessions
- View and manage Symmetrix system migration sessions

- View and manage replication groups and pools
- View and manage SAN LUNs, Open Replicator, RecoverPoint™, and Federate Live Migration

The **Data Protection** menu is only available when you select an individual storage system from the system selector.

Performance The **Performance** menu provides access to the following operations:

- 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 storage 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 real-time data
- Create, schedule, execute, and export data queries
- Configure the performance database and perform backups and restores

Databases The **Database** menu provides access to the following operations:

- View storage system-visible databases
- Launch Database Storage Analyzer for optimizing database performance

System The **System** menu provides access to the following operations:

- View system dashboards, from which you can replace failed spare drives, and perform Symmetrix health checks
- View and filter records from the system audit log
- View, acknowledge, and delete alerts for the specified Symmetrix system
- View and manage active jobs
- View storage system hardware
- View iSCSI dashboard
- View File dashboard
- Licensing

The **System** menu is only available when you select an individual storage system from the system selector.

Support The **Support** menu provides access to the online help.

Help:

Displays content-sensitive help for current view or dashboard.

Contents:

The main section of the interface where you will perform the majority of your work. This can contain a dashboard; a list view, as shown above; a details view; or a properties view. In either case, this section usually contains a series of controls (buttons) that allow you to act on the displayed content. For example, a volume list may contain Create and Delete buttons. For operations that you can perform on multiple objects at a time (for example, deleting storage groups), Shift and Ctrl keys are supported to allow you to select multiple sequential, and non-sequential objects, respectively.

Status bar:

Provides the following information:

- Link to the **Alerts New** list view for the storage system on which you are currently working. The maximum number of alerts Unisphere for VMAX will display is 10,000 per server. Once this threshold is reached, Unisphere for VMAX will delete the oldest alert for each subsequent alert it receives.
- Link to the **Jobs List** for the selected storage system. This only appears after you have selected a storage system.
- The date and time the current contents (view, dashboard, etc.) was updated.
- when clicked refreshes data in Unisphere console window. Unisphere for VMAX refreshes all its data from the database.
- when clicked refreshes storage system data. Unisphere for VMAX refreshes storage system data from the database. This operation will not discover new storage systems, only refresh data from existing systems.

Using the Home Dashboard

The **Home Dashboard** provides a single place from which you can quickly determine how your storage systems are performing. The **Home Dashboard** opens by default when you first open the Unisphere for VMAX console. To return to it from another view, click in the menu bar.



The **Home Dashboard** displays the following, depending on your storage operating environment:

Note Storage systems appear grayed-out when they are in a disabled state. This usually occurs when the Unisphere for VMAX server managing the storage system is in the process of starting up.

- System nice name, or if not defined, system ID number. Unlicensed systems are indicated with . To license the system, click the icon to open the **Entitlements** page from which you can license the storage system.
- **Model** number
- **Connection** (Local or Remote)
- **Enginuity/HYPERMAX OS** version. Enginuity/HYPERMAX OS versions below the minimum suggested target version display a warning icon . In which case, you should contact your service provider to schedule an upgrade.
- Alert summary icon that displays the alert status and count. Alert status is indicated by icon color:

  One or more fatal alerts

  One or more critical alerts, with none higher

  One or more warning alerts, with none higher

  One or more informational alerts, with none higher

No alerts

To view details on the system's alerts, click the alert count (#) to open the system's alert list view.

- **Capacity** — Graphic representation of the system's capacity (used = blue and free = grey). For storage systems running HYPERMAX OS 5977 or higher, only virtual capacity is displayed. For systems running lower OS versions, both physical and virtual capacities are displayed. Click the **Capacity** link to open the storage system's **Dashboard** (or the SRP dashboard for systems running HYPERMAX OS 5977 or higher). Data only displays in this field for local storage systems.
- **SLO Compliance** — How well the storage system's workload is complying with the overall Service Level Objective (SLO). This field displays for storage systems running HYPERMAX OS 5977 or higher and registered with the Performance component. Data only displays in this field for local storage systems. Possible values are:

  29 Stable — Number of storage groups performing within the SLO targets. indicates that there are no storage groups performing within the SLO targets.

  10 Marginal — Number of storage groups performing below SLO targets. indicates that there are no storage groups performing below SLO targets.

  1 Critical — Number of storage groups performing *well* below SLO targets. indicates that there are no storage group performing well below SLO targets.

  To view the **Storage Groups Dashboard**, click the **SLO Compliance** link. To view the available Service Level Objectives, click .
- **SG Response Time** — Graphic representation of how well the storage system's storage groups are performing. This is determined by analyzing the number of storage groups with more than 10 I/Os per second over the last hour, and then displaying the appropriate icon (green (Stable), yellow (Marginal), red (Critical) disks, as shown earlier for SLO compliance). This field displays for storage systems running Enginuity 5874 - 5876 and registered with the Performance component.
- **System Utilization** — Graphic representation of the distribution of % busy for the storage system's front-end, back-end, and cache partitions. To calculate the % busy, the system runs through the following categories and metrics over four hours to determine if a threshold has been crossed:
  - **Front End** = The following summed together:

    Front end directors: % Busy, % Busy Port 0, % Busy Port 1 and

    RDF directors: % Busy, % Busy Port 0
  - **Cache** = % WP Utilization
  - **Back End** = DA and DX: % Busy, % Busy Port 0, % Busy Port 1

  Results are displayed in bar charts of the following colors:

  Green — Indicates the percent of checks that returned values below the first threshold.

  Yellow — Indicates the percent of checks that returned values between the first and second thresholds.

  Red — Indicates the percent of check that returned values above the second threshold.

Using the system selector



The system selector, located at the second left of the menu bar, displays a list of all managed storage systems. It allows you to toggle the interface between two modes, **All Symmetrix** and individual systems. The mode you select will determine the menus/controls available to you (for example, the menus under are only available when you select **All Symmetrix**) and provide the context for subsequent operations.

For example, to create a host on SYMID001234:

1. Select **SYMID001234** from the selector.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Click **Create** to open the **Create Host** dialog box.

Similarly, selections you make 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 SYMID005678 from the selector, the view changes to list the thin pools for SYMID005678.

In addition, the system selector includes the **Discover Symmetrix** option that enables you to retrieve information on your storage environment. For more information, refer to Discovering storage systems (page 30).

Discovering storage systems

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

To discover a storage system:



1. Click the system selector, located at the far left of the menu 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 in the title bar and select **Export** to open the **Export Wizard**.
2. Expand **Application Container** 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**, or as **Image and Data Table** 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 status bar. Unisphere for VMAX refreshes all its data from the database.

Refreshing storage system information

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

Setting preferences

1. Click in the title bar and select **Settings** 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** — This feature is not supported.
   - **Optimize for Remote Connection** — This feature is not supported.
3. Click **OK**.

Searching for storage objects

To search for objects (storage groups, hosts, initiators) across all manage storage systems:

1. Click in the title bar.
2. Select the type of object (**Storage Group**, **Host**, **Initiator**).
3. Depending on the object you are looking for, type the following:
   - **Storage Group** — Type all or part of the storage group name.
   - **Host** — Type all or part of the host name.
   - **Initiator** — Type all or part of the initiator name.
   - Select **All Symmetrix**or a specific storage system identifier.
4. Click **Find**.

   Results include the object **Name**, the **Object Type**, and the associated storage system (**Symmetrix ID**).
5. To view object details, click the object name to open its **Details** view.
6. Click **Clear** to clear the results of the search.

Filtering lists of objects

Unisphere for VMAX includes two types of filtering tools to help you quickly locate specific objects in lists, wizards, and dialog boxes:

- **Simple** — Allows you search for specific objects by name. For example, the name of a storage group.
- **Advanced dialog** — Allows you to filter lists by attributes. For example, capacity, reservations, etc.

Note The system **Audit Log** and **Alerts** list views include their own custom filters. For instructions on using these filters, refer to Filtering audit log records on page 820 and Filtering alerts on page 46, respectively. In addition, the **Volumes** view for storage systems running HYPERMAX OS 5977 or higher includes another form of advanced filter. For instructions on using it, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208.

To use the simple filter:

1. Click the filter icon and click **Show Filter** 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** — 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. (This option is only available on storage systems running Enginuity 5773 - 5876.)
     - **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.

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. This option is only available on storage systems running Enginuity 5875.139.93 - 5876.
  - **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. This option is only available on storage systems running Enginuity 5671 - 5876.
  - **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.
- **Data Domain** — Specifies whether to include/exclude Data Domain volumes. This filter is only available on storage systems running HYPERMAX OS 5977 or higher.

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.

Federated Tiered Storage

Note For storage systems running HYPERMAX OS 5977 or higher, this option is only available when using ProtectPoint.

- **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 in the title bar and select **Logout**.

Getting help

Clicking in the title bar and selecting **Help** 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 3
Administration

This chapter contains the following:

Setting system preferences 38
Displaying login messages 41
Database backup 42
Alert settings 44
Security 85
Link and launch 101

Setting system preferences

Before you begin:

Only a user with Administrator permission can set preferences.

To set system preferences:

1. From the system selector, select **All Symmetrix**.
2. Select **Home › Administration** to open the **Administration** page.
3. Click **Preferences** to open the **Preferences** page.
4. Modify any number of the following preferences:
   - **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 storage system volume reservation preferences. Set the following parameters:

     Note 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 starts** — 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 Pair 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. Possible values are:
  - **No Copy No Diff** — Create a nondifferential (full) copy session without a full background copy.
  - **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 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:
  - **Best Effort** — Specifies to satisfy the request from existing volumes, and then create the volumes necessary to meet any shortfall.
  - **Find Existing** — Specifies to select from existing volumes.
  - **Create New** — Specifies to create new volumes.
- **Clone Target** — Select the default target volume.
- **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. Enabling this option limits the eligible volumes. This option is only available on storage 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 storage 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.
- **Protection Setup Wizard SRDF Communication Protocol** — Select the default SRDF communication protocol, **Fibre Channel** or **GigE**.
- **Protection Setup Wizard SRDF Number of Ports** — Select the default number of ports to use with SRDF.

◆ **SRDF/S Response Time (ms)** — Specify the threshold values used to determine the status (color) of the **Remote Replication** tile in the **Protection Dashboard**. For example, with the default values (Green = 5 and Yellow = 10) set, the following applies:

- If the response time of the majority of remotely protected device groups is less than 5, the status will display green.
- If the response time of the majority of remotely protected device groups is less than 10, the storage group willd yellow.
- If the response time of the majority of remotely protected device groups is more than 10, the storage group is considered red.

◆ **Client Session** — Type a message to display to Unisphere for VMAX users during login. For example, you may want to notify logging in users about a software upgrade. Messages can be up to 240 characters/6 lines.

5. 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. Select **Home › Administration › Preferences** to open the **Preferences** page.
3. In the **Client Session** panel, type the message you want to display. Messages can be up to 240 characters/6 lines.

Database backup

Backing up the database server

This procedure explains how to backup all the data currently on the database server, including Database Storage Analyzer, Workload Planner, performance, and infrastructure data. Database backups will enable you to recover from system crashes.

You can only restore the database to the same Unisphere for VMAX version and same operating system. For example, a V8.0.1 database on Windows, can only be restored to a V8.0.1 on Windows.

Before you begin:

To perform this operation, you must be an Administrator.

To backup the database server:

1. From the system selector, select **All Symmetrix**.
2. Select **Home › Administration › Databases** to open the **Databases** list view.
3. Click **Backup** to open the Database Backup dialog box.
4. In **File Name**, type a description of the backup. Note that the final file name will consist of a time stamp and your custom description.
5. Click **OK**.

Viewing database backups

This procedure explains how to view database backups.

Before you begin:

To perform this operation, you must be a Monitor.

To view database backups:

1. From the system selector, select **All Symmetrix**.
2. Select **Home › Administration › Databases** to open the **Databases** list view.

   The following properties display:

   - **Backup Name** — Name of the backup in the form *TimeStamp_CustomName*.
   - **Status** — Status of the backup.
   - **Start Time** — Time the backup started.
   - **End Time** — Time the backup ended.
   - **Message** — Message related to the backup.

   The following controls are available:

   - **Backup** — Backing up the database server on the previous page
   - **Delete** — Deleting database backups below

Deleting database backups

Before you begin:

To perform this operation, you must be an Administrator.

To delete database backups:

1. From the system selector, select **All Symmetrix**.
2. Select **Home › Administration › Databases** to open the **Databases** list view.
3. Select one or more backups and click **Delete**.
4. Click **OK** in the confirmation message.

Alert settings

You can configure Unisphere for VMAX to monitor storage 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 on page 83.

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

Alerts

Viewing alerts

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

Before you begin:

- For alert (event) descriptions, refer to Alert descriptions on page 49.
- 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 83.
- The maximum number of alerts Unisphere for VMAX will display is 10,000. Once this threshold is reached, Unisphere for VMAX will delete the oldest alert for each subsequent alert it receives.

To view alerts:

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

   For a particular storage system:

   a. Select the storage system.

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

   For all visible storage systems:

   a. Select **All Symmetrix**.

   b. Select **Home** › **Alerts** to open the **Alert** list view.

2. *Optional*: Use the alert filter to view a subset of the listed alerts. For more information on the Alert Filter, refer to Filtering alerts (page 46).

In both cases, the following properties display:

- **State** — State of the alert. Possible values are New or Acknowledged.
- **Severity** — Alert's severity. Possible values are:
  - Fatal
  - Critical
  - Warning — The following events map to this severity:
    - The component is in a degraded state of operation.
    - The storage array is no longer present (during certain operations).
    - The component is in an unknown state.
    - The component is (where possible) in a write-disabled state.
  - Information — The component is no longer present (during certain operations).
  - Normal — The component is now (back) in a normal state of operation.
- **Type** — Type of alert. Possible values are Array, Performance, Server, System, and File.
- **Symmetrix** — Storage 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** — Component to which the alert is related. 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.

The following controls are available:

- **View Details** — See Viewing alert details on the next page.
- **Acknowledge** — See Acknowledging alerts on the next page.
- **Delete** — See Deleting alerts on page 47.

Filtering alerts

The **Alerts** list view (**System › Alerts**, or **Home › Alerts**) includes a filter tool for narrowing the listed alerts to only those that meet the specified criteria:

- **Start Date** — Filters the list for alerts created on the specified date/time.
- **End Date** — Filters the list for alerts acknowledged on the specified date/time.
- **State** — Filters the list for alerts with the specified state.
- **Severity** — Filters the list for alerts with the specified severity.
- **Type** — Filters the list for alerts with the specified type.
- **Object** — Filters the list for alerts for the specified object.
- **Description** — Filters the list for alerts with the specified description.

Acknowledging alerts

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

Viewing alert details

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

   This view contains two panels, Properties and Related Objects.

Properties panel

The following properties display:

- **Alert ID** — Unique number assigned by Unisphere for VMAX.
- **State** — State of the alert. Possible values are new or acknowledged.
- **Severity** — Alert's severity. Possible values are:
  - Fatal
  - Critical
  - Warning
  - Information
  - Normal
- **Type** — Type of alert. Possible values are Array, Performance, and System.
- **Symmetrix** — ID of the storage system generating the alert.

- **Object** — Object to which the alert is related. For more information, click the object to open its details view.
- **Created** — Date/time the alert was created.
- **Description** — Description of the alert.

The following controls are available:

- **Acknowledge** — See Acknowledging alerts on the previous page.
- **Delete** — See Deleting alerts below.

Objects panel

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

Deleting alerts

This procedure explains how to delete System and Array alerts:

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

Configuring alert notifications

This procedure explains how to configure Unisphere for VMAX to notify you when a storage system generates an alert.

Before you begin:

- To perform this operation, you must be an Administrator or StorageAdmin.
- Unisphere for VMAX employs the following throttling algorithms to prevent alert flurries from straining the system.
  - Storage system Event Throttling — When a storage 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-storage sysevent exceeds a set threshold, the alert infrastructure will ignore subsequent alerts from the source.
- Alert notifications for runtime checks will not be delivered. For more information on runtime checks, refer to Runtime checks on page 83.

To enable alert notifications:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › Notifications** to open the **Notifications** page.
3. In the **Types** panel, click **Enable** next to the method you want to use to deliver the notifications. Possible methods are:
   - **Syslog** — Forwards alert notifications to a remote syslog server.
   - **Email** — Forwards alert notifications to an email address.
   - **SNMP** — Forwards alert notifications to a remote SNMP listener.
4. Depending on the method you selected, do the following:
   Syslog:

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

   Email:

   a. Type the email address of the email server (**Sender Email**).
   b. Type the IP address of the email server (**IP Address/Host**).
   c. If the port of your email server differs from the default (25), type a new value (**Server Port**).
   d. Type the email address to which notifications will be sent (**Recipient Email**l).
   e. *Optional*: To send a test message to the email server, click **Test**.
   f. Click **OK**.

SNMP:

a. *Optional*: Click **Add SNMP** to add another SNMP server and click **OK** to confirm.

a. *Optional*: Click X to remove a SNMP server and click **OK** to confirm.

b. Type IP address of the SNMP server (**IP Address/Host**).

c. Type the port on the SNMP server (**Server Port**).

d. *Optional*: To send a test SNMP trap, click **Test**.

e. Click **OK**.

1. In the **Alerts** panel, do the following for each storage system from which you want to receive notifications:

   a. Select the **System Level** and **Performance Level** severites in which you are interested.

   b. Type the email addresses to which the notifications should be sent.

   c. To clear your changes, click **Discard**.

   d. Once satisfied, click **Save**.

To disable alert notifications:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › Notifications** to open the **Notifications** page.
3. Click **Disable** next to method you no longer want to use.

Alert descriptions

Unisphere for VMAX reports the following alerts:

Message: event daemon restarted; events may have been lost

Category: N/A

Object: N/A

Severity: Warning

Notes: Generated when the event daemon is restarted after a crash.

Message: event daemon communications problem; events may have been lost

Category: N/A

Object: N/A

Severity: Warning

Notes: Generated when the event daemon encounters a communication problem attempting to send events back to a client.

Message: Event Queue overflow; events may have been lost

Category: N/A

Object: N/A

Severity: Warning

Notes: Generated when one of the internal Event Queues (within a client process or event daemon) overflows and events are discarded.

Message: Device state has changed to Online [Degraded, RG-Mbr-Rebuild-or-Copy

Category: Status

Object: Device number

Severity:

If Online = Normal

If Online = Degraded Warning

If Offline = Critical

If Not Present = Info

Notes:

'Not Present' means that the device could not be seen by Solutions Enabler.

'Online' means that the device service state is normal.

'Online [Degraded]' means one or more of the device's mirrors are in a Not-Ready state.

'Offline' means that the device service state is failed.

Message: Array state has changed to Not Present | Unknown | Online | Online Degraded | Offline

Category: Status

Object: N/A

Severity:

If Online = Normal

If Online Degraded = Warning

If Offline = Critical

If Not Present = Warning or Critical (depending on the situation)

If Unknown = Warning or Critical (depending on the situation)

Notes:

This event reflects the overall state of the array - including its Disks, Directors, Ports.

'Not Present' means that the array couldn't be seen.

'Online' means that the array is operational.

'Online [Degraded]' means that:

One or more Ports are in an Offline or Write-Disabled state.

One or more Directors are in an Offline or Dead state.

Device events [1200] events are also enabled and one or more device is in a Not-Ready state.

Array sub-component events [1404] are also enabled and one or more are in a failed (Offline) state (Fans, Power Supplies, LCCs, MIBEs, Enclosures, etc.).

'Unknown' means that there was a problem communicating with the array.

Message: Director state has changed to Not Present | Online | Online Degrade] | Offline | Failed

Category: Status

Object: Director identifier

Severity:

If Online = Normal

If Online Degraded = Warning

If Offline = Fatal

If Not Present = Info

Notes:

'Not Present' means the director was not seen by Unisphere for VMAX.

'Online' means that the director status is Online.

'Online [Degraded]' means that one or more of the director's ports were in an Offline or Write-Disabled state.

'Offline' means that the director status is Offline.

'Failed' means that the director status is Dead.

Message: Port state has changed to Not Present | Unknown | Online | Write Disabled | Offline

Category: Status

Object: Port identifier

Severity:

If Online = Normal

If Offline = Critical

If Write Disabled = Warning

If Unknown = Warning

If Not Present = Info

Notes:

'Not Present' means the port was not seen.

'Online' means a port status of On.

'Offline' means a port status of Off.

'Write Disabled' means a port status of Write-Disabled.

Message: Disk state is now *State* (was: *State*).

Where *State* can be:

Online

Offline

Online Spare

Offline Spare

Online Degraded

Not Present

Category: Status

Object: Spindle ID (Disk identifier is supported for internal disks only)

For example:

Disk=123

or

Disk=16B:C2 (for Director 16B, DA Interface C, SCSI ID/Target 2) (internal disks only)

Severity:

If Online = Normal

If Online = Spare Normal

If Online = Degraded Warning

If Offline = Warning

If Offline = Spare Warning

If Not Present = Info

Notes:

'Not Present' means that the disk could not be seen by Unisphere for VMAX.

'Online' means that one or more of the disk's Hypers are in a Ready state.

'Online Spare' means that the disk is a Spare and one or more of the disk’s Hypers are in a Ready state.

'Online [Degraded]' means that the disk can only be reached via a single array DS controller. This disk state is for external disk only and supported with Enginuity 5876 and later.

'Offline' means that all of the disk's Hypers are in a Not-Ready state.

'Offline Spare' means that the disk is a Spare and all of the disk's Hypers are in a Not-Ready state.

Message: Device configuration has changed

Category: Status

Object: Volume number

Severity: Info

Notes: The following aspects of a device's configuration is considered by this event:

The base device configuration.

The meta configuration of the device (META_HEAD, META_MEMBER).

The bound-vs-bound state of a TDEV (bound vs unbound).

Whether a dynamic spare disk is invoked for the device.

The RDF mode of the device (of either leg for Concurrent SRDF).

The data pool bound to by a TDEV changes. This reflects a device being bound, unbound or re-bound to a different pool, and is also triggered when the name of the pool changes.

Message: Snap Savedev Pool state has changed to Not Present | Online | Online Degraded | Offline
Message: SRDF/A DSE Pool state has changed to Not Present | Online | Online Degraded | Offline
Message: Data Pool state has changed to Not Present | Online | Online Degraded | Offline

Category: Status

Object: Pool name

Severity:

If Online = Normal

If Online Degraded = Warning

If Offline = Critical

If Not Present = Info

Notes:

'Not Present' means that the pool no longer exists.

'Online' means that the pool is in an enabled.'

'Online [Degraded]' means that the pool is in a mixed state.

'Offline' means that the pool is in a disabled state.

Message: Snap Savedev Pool configuration has changed
Message: SRDF/A DSE Pool configuration has changed
Message: Data Pool configuration has changed

Category: Status

Object: Pool name

Severity: Info

Notes:A pool's configuration changes if:

The set of Enabled devices in the pool changes.

The total size (free + used) of all the Enabled devices in the pool changes.

Message: Snap Savedev Pool utilization is now ‹NN› percent
Message: SRDF/A DSE Pool utilization is now ‹NN› percent
Message: Data Pool utilization is now ‹NN› percent

This is a threshold event that reflects the amount (as a percentage) of used space within a pool.

Category: Status

Object: Pool name

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is 100% — Fatal

If value is >= 80% — Critical

If value is >= 70% — Major

If value is >= 65% — Minor

If value is >= 60% — Warning

Otherwise — Nornal

Notes:

Used space size is determined by calls to SymPoolShow().

Events are only delivered at multiples of 5% ... for <NN> equal to 5%, 10%, 15%, ... , 75%, 80%, 85%, 90%, 95% and 100%.

Threshold events are only delivered when the severity, as determined by threshold values, changes.

Message: Symmetrix External Lock has been acquired
Message: Symmetrix External Lock has been released

Indicates that one of the monitored Symmetrix External Locks (SELs) has been acquired or released.

Category: Status

Object: Symmetrix lock number

Severity: Info

Notes: At this time, only SEL #15 (used by Config Change) is monitored.

Indicates that one of the monitored Symmetrix External Locks (SELs) has been released.

Category: Status

Object: Symmetrix lock number

Severity: Info

Notes: At this time, only SEL #15 (used by Config Change) is monitored.

Message: Disk is no longer a Spare
Message: Disk is now a Spare
Message: Disk now an invoked Spare
Message: Spare has been invoked against a failed disk
Message: Spare is no longer invoked against a failed disk

Indicates that a disk has started or stopped acting as a spare.

Category: Status

Object: Disk identifier of the spare.

Severity:

For 5x74 and newer arrays:

Normal

For older arrays:

If invoked = Warning

If no longer invoked = Normal

Notes: With Permanent Sparing on newer arrays, a failing disk and a spare will exchange roles. The failed disk will end up as a failed spare, and the spare will end up as a normal disk. The "Disk is now an invoked Spare" event will rarely if ever be delivered.

Message: Number of available disk spares is ‹NN›

This is a Threshold event that reflects the number of available Spare Disks on the storage array.

Category: N/A

Object: N/A

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is 0 — Critical

If value is 1 — Major

If value is 2 — Warning

Otherwise — Info

Notes: Threshold events are only delivered when the severity, as determined by threshold values, changes.

Message: Number of available disk spares is ‹NN›

This is a Threshold event that reflects the number of available Spare Disks on the storage array.

Category: N/A

Object: N/A

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

TBD

Notes: Threshold events are only delivered when the severity, as determined by threshold values, changes.

Messsage: Thin Device is now ‹NN› percent allocated.

This is a Threshold event that reflects the amount (as a percentage) of a Thin Device that is backed by space in a Data Pool.

Category: N/A

Object: Device number

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is 100% — Fatal

If value is >= 80% — Critical

If value is >= 70% — Major

If value is >= 65% — Minor

If value is >= 60% — Warning

Otherwise — Normal

Notes: Events are only delivered at multiples of 5% ... for ‹NN› equal to 5%, 10%, 15%, ... , 75%, 80%, 85%, 90%, 95% and 100%.

Threshold events are only delivered when the severity, as determined by threshold values, changes.

Message: Thin Device is now ‹NN› percent allocated

This is a Threshold event that reflects the amount (as a percentage) of a Thin Device that has been written to.

Category: N/A

Object: Device number

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is 100% — Fatal

If value is >= 80% — Critical

If value is >= 70% — Major

If value is >= 65% — Minor

If value is >= 60% — Warning

Otherwise — Normal

Notes: Events are only delivered at multiples of 5% ... for ‹NN› equal to 5%, 10%, 15%, ... , 75%, 80%, 85%, 90%, 95% and 100%.

Threshold events are only delivered when the severity, as determined by threshold values, changes.

Message: Port configuration has changed

This is a Threshold event that reflects the amount (as a percentage) of a Thin Device that has been written to.

Category: Status

Object: Port

Severity: Info

Notes: The only aspects of a port's configuration that are considered the following flags from the FA port flags:

The _VCM_ENABLED flag

The _VOL_SET_ADDR (VSA) flag

Message: Snap Savedev Pool device state has changed
Message: SRDF/A DSE Pool device state has changed
Message: Data Pool device state has changed

Indicates that the Deferred Service replacement threshold indicator for a storage array has changed. This change can be in either direction – from not-exceeded to exceeded ... or from exceeded to not-exceeded.

Category: Status

Object: N/A

Severity:

If the replacement threshold has been exceeded = Warning

Otherwise = Info

Notes: This event will only be generated if Deferred Service is enabled for the storage array.

Message: The device configuration checksum has changed

Indicates that the configuration of a device has changed. The implementation makes use of a checksum maintained by the event daemonover a device's core configuration data. An event is generated when this checksum changes.

Category: Status

Object: Volume number

Severity: Info

Notes: N/A

Message: The migrate operation is complete: success
Message: The migrate operation is complete: timed out
Message: The migrate operation is complete: terminated
Message: The migrate operation is complete: failed

Indicates that a VLUN migration has completed or failed.

Category: Status

Object: Migrate session name

Severity:

If success = Info

If terminated = Info

If timed out = Warning

If failed = Major

Notes: This is only generated for explicitly initiated VLUN migrations – not movements being performed by FAST.

Message: Thin Pool rebalancing operation is complete: success
Message: Thin Pool rebalancing operation is complete: timed out
Message: Thin Pool rebalancing operation is complete: terminated
Message: Thin Pool rebalancing operation is complete: failed

Indicates that a Thin Pool rebalancing activity has completed.

Category: Status

Object: Data pool name

Severity:

If success = Info

If terminated = Info

If timed out = Warning

If failed = Major

Notes: This event is only supported for Symmetrix arrays running Enginuity 5875.

Message: Array configuration has changed.

Indicates that some change has been made to the configuration of the storage array.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: Device Masking database has changed

Indicates that some change have been made to the device masking database on the storage array.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: Access Control definitions have changed

Indicates that some change has been made to the Access Control [symacl] database on the storage array.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: Dynamic RDF operation performed on device

Indicates that a dynamic RDF operation has been performed on some device.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: Snap session created, activated or deleted

Indicates that a snap / clone session has been created, activated or deleted.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: BCV device pairing has changed

Indicates that the BCV pairing for some device has changed.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: HPUX device identifier has changed

Indicates that the HPUX device identifier for some device has been changed.

Category: Status

Object:

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: Device Name has changed

Indicates that the device name for some device has been changed.

Category: Status

Object: N/A

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: Device Nice Name has changed

Indicates that the device name for some device has been changed.

Category: Status

Object: N/A

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Message: OpenVMS device identifier has changed

Indicates that the OpenVMS device identifier for some device has been changed.

Category: Status

Object: N/A

Severity: Info

Notes: This event is derived from one of the QuickConfig indication maintained on the storage array.

Device Reservations data has changed

Indicates that the Device Reservation state for some device on the storage array has changed.

Category: Status

Object: N/A

Severity: Info

Notes: This event requires checking for modifications to file(s) within SFS.

Message: Time since last SRDFA cycle switch exceeds minimum cycle time by ‹NN› seconds

This is a threshold event that indicates the amount (in seconds) by which an SRDFA group's cycle time exceeds the minimum that is configured.

Category: N/A

Object: SRDF group

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is ›= 5 — Warning

Otherwise — Info

Notes: This is determined by calling SymReplicationGet() and examining the time_since_last_switch and duration_of_last_cycle quantities for Active, R1, non-MSC sessions.

The event value corresponds to the number of seconds that the larger of these two is beyond the configured min_cycle_time. If the time(s) are less than min_cycle_time (everything normal), the event value is 0. To protect against rounding problems, the test is actually against min_cycle_time+1. If the times are less than min_cycle_time+1, the event value will be 0. Therefore, possible event values are: 0, 2, 3, 4, 5, etc.

For example, assuming a min_cycle_time of 10:

Time since last switch	Event value
9	0
10	0
11	0
13	3

Message: SRDFA cycles now using ‹NN› percent of the cache available for it

Category: N/A

Object: N/A

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is ›= 90 — Warning

Otherwise — Info

Notes: This is a Threshold event that indicates the percentage of cache that is available for SRDFA use that is actually holding SRDFA Write Pending data.

This is determined by calling SymReplicationGet() and summing the active_cycle_size and inactive_cycle_size values for all active R1 or R2 sessions. The maximum available cache is computed in the usual manner:

```
if ((max_host_throttle == 0) and
(rdfa_max_cache_usage › 0) and
(rdfa_max_cache_usage ‹ 100))
max_avail = (max_wr_pend_slots * rdfa_max_cache_usage) / 100
else
max_avail = max_wr_pend_slots
```

The event value is the sun of the active and inactive cycle sizes expressed as a percentage of this max avail cache size.

warning: Exercise caution when assigning significance to this event. The fact that an amount of cache is available for SRDFA to use (max_avail above) doesn't mean that it is guaranteed to be available for its use. There are other sources of Write Pending data that can use up this space as well - leaving it unavailable for SRDFA's use.

Message: Write Pending data is now using ‹NN› percent of the cache

Category: N/A

Object: N/A

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is >= 90 — Warning

Otherwise — Info

Notes: This is a Threshold event that indicates the percentage of Symmetrix Cache that is holding Write Pending data.

Message: Component state has changed to Online | Online Degraded | Offline | Unknown

Category: N/A

Object:

Power=xxx

Fan=xxx

LCC=xxx

Enclosure=xxx

MM=xxx

IOMC=xxx

Dir=xxx

Severity:

If Online = Normal

If Online = Degraded Warning

If Offline = Major

If Unknown = Warning

Notes: Indicates a change in environmental status for one of the following types of sub-components within the VMAX array:

Fans	[Fan]
Power Supplies	[Power]
Link Control Cards	[LCC]
Management Modules	[MM]
IO Module Carriers	[IOMC]
Directors (for environmental alerts)	[Dir]
Enclosures or Matrix Interface Board Enclosures	[Enclosure]

Online' means that the component is a Normal or Degraded state.

'Online [Degraded]' means that the component is in a degraded state.

'Offline' means that the component is in a Failed state.

The format of the specific component name ('xxx' above) may differ depending on the VMAX model. Some examples you might encounter are:

SB-1/Fan-A	Fan in System Bay
SB-1/ENC-1	Enclosure within System Bay
SB-1/ENC-1/Fan-A	Fan in Enclosure-Slot within System Bay
SB-1/MIBE-L-2A	MIBE within System Bay
SB-1/MIBE-L-2A/PS-A	Power Supply in MIBE within System Bay

Message: DSE Spillover has been occurring on the RDF group for ‹N› minutes
Message: DSE Spillover is no longer occurring on the RDF group

Category: N/A

Object: The SRDF group

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is ›= 30 (minutes) — Warning

Otherwise — Normal

Notes: Threshold events are only delivered when the severity, as determined by threshold values, changes.

Message: Disk Group has changed
Message: Disk Group has been deleted
Message: Disk Group has been created

Category: N/A

Object: The Symmetrix disk group number (decimal)

Severity: Info

Notes: This event is only supported on VMAX arrays running Enginuity 5876 and HYPERMAX OS 5977.

Message: Disk has spare coverage
Message: Disk no longer has spare coverage

Category: N/A

Object: Disk identifier

Severity: Info

Notes: This event is only supported on VMAX arrays running Enginuity 5876 and HYPERMAX OS 5977.

Message: The spare capacity percentage for diskgroup on engine %d has changed

Category: N/A

Object: Symmetrix Disk Group number (decimal)

Severity: Info

Notes: This event is only supported on VMAX arrays running Enginuity 5876 and HYPERMAX OS 5977.

Message: Cache Partitioning configuration has changed

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Dynamic Mapping configuration for a device has changed

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Meta configuration for a device has changed

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Initiator Group has changed

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Storage Group has changed

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Masking view has changed

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Feature registration DB has changed

Indicates that a change has been made to the Feature Registration DataBase on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: FAST tiers have changed

Indicates that a change has been made to the FAST (Fully Automated Storage Tiering) Tiers on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: FAST policies have changed

Indicates that a change has been made to the FAST (Fully Automated Storage Tiering) Policies on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: FAST associations have changed

Indicates that a change has been made to the FAST (Fully Automated Storage Tiering) Associations on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Optimizer/FAST time windows have changed

Indicates that a change has been made to the FAST (Fully Automated Storage Tiering) time windows on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Optimizer/FAST control parameters have changed

Indicates that a change has been made to the FAST (Fully Automated Storage Tiering) control parameters on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array.

Message: Storage group configuration has changed

Indicates that a change has been made to a storage group on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: Initiator group configuration has changed

Indicates that a change has been made to an initiator group on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: Port group configuration has changed

Indicates that a change has been made to a port group on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: Masking view configuration has changed

Indicates that a change has been made to a masking view configuration on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: Storage group scope has changed

Indicates that a change has been made to a storage group scope on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: Initiator group scope has changed

Indicates that a change has been made to an initiator group scope on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: Masking view scope has changed

Indicates that a change has been made to a masking view scope on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is derived from one of the QuickConfig indication maintained on the VMAX array. This event is only supported with HYPERMAX OS 5977.

Message: User Authorization rules have changed

Indicates that a change has been made to the User Authorization [symauth] database on the VMAX array.

Category: N/A

Object: N/A

Severity: Info

Notes: This is determined by checking for modifications to the User Authorization file stored in SFS.

Message: Audit log is at ‹NN› percent of capacity (before wrapping)

This is a threshold event that tracks as a percentage the amount of data in an array Audit Log - how close the log is to its wrapping point where existing entries begin to be over-written.

Category: N/A

Object: N/A

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is ›= 80% — Warning

Otherwise — Normal

Notes: What is actually reported is the position of the write pointer within the Audit Log as a percentage: 0% for the beginning, 100% for the end.

This event is intended to be used as an indication that a backup of the Audit Log is needed - if appropriate.

Message: ‹‹ The actual message from the Audit Record ››

Indicates that a security-related record was written to the array Audit Log.

Category: N/A

Object: N/A

Severity: Info

Notes: This event is delivered when audit records with an Audit Class of SECURITY are detected in the Audit Log.

The audit message is a free-form string that may span multiple lines (containing multiple new line characters).

Message: ‹‹ The actual message from the Audit Record ››

Indicates that a Security alert was written to the array Audit Log.

Category: N/A

Object: N/A

Severity: Info

Notes:

- This event is delivered when audit records with an Audit Class of SECURITY are detected in the Audit Log.
- Access Control failures (host based access control, symacl).
- User Authorization failures (user based access control, symauth).
- SymmWin / SSC Logon failures.
- SymmWin Logins.
- iSCSI authorization failures.
- The audit message is a free-form string that may span multiple lines (containing multiple new line characters).

Message: ‹‹ The actual message from the Audit Record ››

Indicates some (any) record written to the array Audit Log.

Category: N/A

Object: N/A

Severity: Info

Notes: The audit message is a free-form string that may span multiple lines (containing multiple new line characters).

Message: Optimizer Swap activity (from Audit Log)

Indicates some Optimizer Swap activity.

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: Optimizer Move activity (from Audit Log)

Indicates some Optimizer Move activity.

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: Optimizer configuration change (from Audit Log)

Indicates some Optimizer configuration change.

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: FAST Controller Swap activity (from Audit Log)

Indicates some FAST Controller activity.

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: FAST Controller Move activity (from Audit Log)

Indicates some FAST Controller Move activity.

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: FAST Controller configuration change (from Audit Log)

Indicates some FAST Controller configuration change.

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: Optimizer Rollback activity (from Audit Log)

Category: Optimizer

Object: N/A

Severity: Info

Notes: This is derived by detecting a record written by the Optimizer to the array Audit Log.

Message: User approval is required for a Config Change plan generated by the Optimizer/FAST Controller

Indicates that user approval of the a swap state is required and user approval is required.

Category: N/A

Object: N/A

Severity: Info

Notes: N/A

Message: The FAST (DP or VP) controller has switched to state: ‹current_state› (was: ‹previous_state›)

Where ‹current_state› and ‹previous_state› can be one of the following possible values:

- Disabled
- Enabled
- Disable in progress
- Enable in progress
- Disable with error
- Degraded

Indicates that the FAST state has changed from a previous state to the current state.

Category: Optimizer

Object: N/A

Severity: Info

Notes: N/A

Message: The Optimizer has switched to a different mode

Indicates that the Optimizer status state has changed.

Category: Optimizer

Object: N/A

Severity: Info

Notes: N/A

Message: The combined allocation in pools has changed

This event checks for allocated capacity change of all associated pools under the same FAST VP policy. And as such, if FAST VP policy is accidentally used, this event will never be generated.

Category: Optimizer

Object: ‹Policy_Name›

Severity: Info

Notes: N/A

Message: FAST Tier (‹tier_name›) is performing as expected (NORMAL).
Message: FAST Tier (‹tier_name›) is performing worse than expected (LOW).

Category: Optimizer

Object: ‹tier_Name›

Severity: Info

Notes: This event is only supported with Enginuity 5876 Q42012 SR and above.

Message: The allocation for SRP has changed to ‹NN› %u percent

Category: Optimizer

Object: ‹SRP_Name›

Severity: Determined by Threshold values. Unless threshold values are supplied with the registration, the following defaults are used to derive an event severity:

If value is ›= 100% — Fatal

If value is ›= 80% — Critical

If value is ›= 70% — Major

If value is ›= 65% — Minor

If value is ›= 60% — Warning

Otherwise — Normal

Notes:

- Events are only delivered at multiples of 5% ... for ‹NN› equal to 5%, 10%, 15%, ... , 75%, 80%, 85%, 90%, 95% and 100%.
- Threshold events are only delivered when the severity, as determined by threshold values, changes.
- This event is only supported with HYPERMAX OS 5977.

Message: Group has changed

Indicates that the composition of a device group (DG) or composite group (CG) has changed.

Category: N/A

Entity: Not set -- set to NULL.

Object: DG or CG

Severity: Info

Notes:

- The Entity name and type (normally a Symmetrix ID) are not provided for this event. When registering to receive the event, there is no need to supply an ID (symid=000194900123) - if one is supplied, it will be ignored.
- If GNS is not enabled, this event indicates that a group definition in the Solutions Enabler DB file on this host has changed.
- If GNS is enabled, this event indicates that a global group definition stored within GNS (on storage arrays) has changed.

Alert policies

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 *Solutions Enabler Installation Guide*.

To enable/disable alert policies:

1. From the **System Selecyetor**, select **All Symmetrix**.
2. Select **Home › 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 policy:

   Enabling:

   a. Select one or more alert policies, and click **Enable**.

   b. To receive a notification when a storage system reports an alert, click **Notify** to open the **Define Notification** dialog box.

   c. Select one or more notification types.

   d. Click **OK** to close the **Define Notification** dialog box.

   Disabling:

   Select one or more alert policies, and click **Disable**.

Viewing alert policies

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › 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** — Storage 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 array-based alerts.
     - **SMAS** for application-based alerts.
     - **File** for eNAS-based alerts.
   - **Policy Name** — Name of the policy. For alert (event) descriptions, refer to the *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 Alert policies on the previous page.
   - **Disable** —See Alert policies on the previous page.
   - **Notify** — See Configuring alert notifications on page 48.

Pool threshold alerts

Managing pool threshold alerts

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.

For information on managing performance thresholds, refer to:

- Performance Thresholds and Alerts on page 631
- Editing a performance threshold and alert on page 632
- Deleting performance thresholds and alerts on page 632

Before you begin:

- Pool utilization thresholds are enabled by default on every VMAX Family system.
- To receive utilization threshold alerts, you must enable alerts on the storage system.
- To receive alert notifications, you must first configure the alert notifications feature.

To manage pool alert thresholds:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › Symmetrix Threshold Alerts** to open the **Symmetrix Threshold Alerts** 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 storage 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. This category only applies to Enginuity 5773 - 5876.
      - **DSE Spill Duration** — Threshold event that reflects how long (in minutes) an SRDF spillover has been occurring. This category only applies to Enginuity 5773 - 5876.
      - **Snap Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within a snap pool. This category only applies to Enginuity 5773 - 5876.
      - **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. This category only applies to Enginuity 5875 - 5876.
      - **Storage Resource Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within an SRP. This category only applies to storage systems running HYPERMAX OS 5977 or higher.

   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 Policy** dialog box.

   b. Specify a new threshold value (percentage of utilization) for any number of the severity levels: **Warning**, **Critical**, and **Fatal**.

   c. Enable (select) or disable (clear) alerts for the threshold.

   d. Click **OK** to close the **Edit Threshold Policy** dialog box.

   Deleting:

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

Viewing pool threshold alerts

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › Symmetrix Threshold Alerts** to open the **Symmetrix Threshold Alerts** list view.

   Use the **Symmetrix Threshold Alerts** list view to display and manage pool threshold alerts.

   The following properties display:

- **Symmetrix ID** — Storage 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 Spill Duration** — 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.
  - **Symmetrix Resource Pool Utilization** — Threshold event that reflects the allocated capacity (as percentage) within an SRP. This category only applies to storage systems running HYPERMAX OS 5977 or higher.
- **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 Pool threshold alerts on page 74.
- **Edit** — See Pool threshold alerts on page 74.
- **Notify** — See Configuring alert notifications on page 48.
- **Delete** — See Pool threshold alerts on page 74.

Performance 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 the Root Cause Analysis view.

To create a threshold:

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Click **Create** to open the **Create Performance Thresholds and Alert** dialog box.
3. Select the **Symmetrix ID**, **Category**, **Instance**, and **Metric**.

   For some categories, when selected, the **Instance** field defaults to **All**.

   If there is a default value for the **Metric**, it automatically displays in the **First Threshold** and **Second Threshold** fields.
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**.
2. Select the severity for the first threshold, the second threshold, or both. The available 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, an alert is initiated.
4. *Optional:* Select any additional configuration options. For some group categories, you can choose to enable for the alert for the individual components of the group, for example, when the **Disk Group** category is selected, you have the option to enable the alert for the disk.
5. 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 that the value has changed from the default.

To edit performance thresholds and alerts

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Navigate to the threshold or alert to be edited by selecting the appropriate category in the **Category** section and expanding the appropriate metric in the **Metric** column.
3. Select the item from the table and click **Edit** to open the **Edit Performance Threshold and Alert** dialog box.
4. Edit the threshold and alert settings, as appropriate.

   For more information about threshold and alert settings, see See "Performance Thresholds and Alerts" on page 631
5. Click **OK**.

Deleting performance thresholds and alerts

Before you begin:

You can delete only custom values. You cannot delete default thresholds.

To delete a performance threshold and alert:

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Navigate to the threshold or alert to be edited by selecting the appropriate category in the **Category** section and expanding the appropriate metric in the **Metric** column.
3. Select one or more rows and click **Delete**.

   A confirmation message displays.
4. In the **Delete Threshold and Alert** confirmation dialog box, click **Delete**.

Viewing thresholds and alerts

To view thresholds and alerts:

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Select the storage system.
3. In the **Category** section, select the category for which you want to view the configured thresholds and alerts.

   The thresholds and alerts configured for that category are displayed, according to metric. By default, the KPI tab displays. To view all metrics, click **All**.

   Any metrics that include a custom threshold or alert are highlighted with a tick mark in the **Custom** column.

The some of the following properties display:

- **Metric** — The storage system identifier.
- **KPI** — Indicates (with a tick mark) if the metric is a KPI.
- **Custom** — Indicates (with a tick mark) whether a custom threshold or alert has been set.
- **First Threshold** — The user-defined first threshold.
- **Second Threshold** — The user-defined second threshold.

The following controls are available:

- **Create** — See Performance Thresholds and Alerts on page 631.
- **Edit** — See Editing a performance threshold and alert on page 632.
- **Delete** — See Deleting performance thresholds and alerts on page 632.
- **Notifications** — See Configuring alert notifications on page 48.

Configuring email notification

You can configure email addresses to which notifications, alerts, and reports are sent. You can configure a single email address for all notification instances, or you can use different email addresses for different notifications on different storage systems.

To set up email notification:

1. Click **Home › Administration › Alert Settings › Notifications**. The **Notifications** window is displayed.
2. In the **Email** section, click **Enable**. The **Configure Email** window is displayed.
3. In the **From** section specify the following details:
   - **Sender Email**
   - **IP Address/Host**
   - **Server Port**
4. In the **Recipient Email** field of the **To** section, specify the address you want to enable for all email notifications.
5. Click **OK**.

To add an email address and apply it to all notification instances:

1. Click **Home › Administration › Alert Settings › Notifications**. The **Notifications** window is displayed.
2. In the **Email** section, click **Enable**. The **Configure Email** window is displayed.
3. In the **Recipient Email** field of the **To** section, specify the address you want to enable for all email notifications.

Note: If you configure a new email address in this way, any previously configured addresses are overwritten.

4. Click **OK**.

To configure email addresses for individual notifications:

1. Click **Home › Administration › Alert Settings › Notifications**. The **Notifications** window is displayed.
2. Click the **Edit** icon for the storage system for which you want to edit the email notification settings. The three individual email fields become editable.
3. Configure the email addresses for individual notifications, as appropriate.
4. *Optional*: Configure the notification levels by clicking the appropriate icons in the **System Level** and **Performance Level** columns.
5. Click **OK**.

SLO compliance alert policies

Creating SLO compliance alerts policies

This procedure explains how to configure Unisphere for VMAX to alert you when the performance of a storage group, relative to its SLO target, changes. Once configured, Unisphere for VMAX will assess the performance of the storage every 30 minutes, and deliver the appropriate alert level.

When assessing the performance for a storage group, Workload Planner calculates its weighted response time for the past 4 hours and for the past 2 weeks, and then compares the two values to the maximum response time associated with its given SLO. If both calculated values fall within (under) the SLO defined response time band, the compliance state is STABLE. If one of them is in compliance and the other is out of compliance, then the compliance state is MARGINAL. If both are out of compliance, then the compliance state is CRITICAL.

The following table details the state changes that will generate an alert and the alert level.

Table 1: SLO compliance rules

State change	Alert generated	Alert level
ANY STATE › NONE	No	
NONE › STABLE	No	
NONE › MARGINAL	Yes	Warning
NONE › CRITICAL	Yes	Critical
STABLE › MARGINAL	Yes	Warning
STABLE › CRITICAL	Yes	Critical
STABLE › STABLE	No	
MARGINAL › STABLE	Yes	Info
MARGINAL › CRITICAL	Yes	Critical
MARGINAL › MARGINAL	No	
CRITICAL › STABLE	Yes	Info
CRITICAL › MARGINAL	Yes	Warning
CRITICAL › CRITICAL	Yes	Critical

Note When a storage group configured for compliance alerts is deleted or renamed, the compliance alerts will need to be deleted manually. For instructions, refer to Deleting SLO compliance alerts policies on page 82.

Before you begin:

- The storage system must be running HYPERMAX OS 5977 or higher and registered for performance stats collection.
- The storage group must:

- Be either a child or standalone. Parent storage groups are not supported.
- Be associated with an SLO other than optimized.
- Contain volumes other than gatekeepers
- Be in a masking view
- Not have a policy currently associated with it

To create SLO compliance alert policies:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › SLO Compliance Alert Policies** to open the **SLO Compliance Alert Policies** list view.
3. Click **Create** to open the **Create Compliance Alert Policy** dialog box.
4. Select the storage system on which the storage groups are located.
5. Select one or more storage groups and click **Add**.
6. *Optional*: By default, SLO compliance policies are configured to generate alerts for all SLO compliance states. To change this default behavior, clear any of the states for which you do not want generate alerts:

   Critical — Storage group performing *well* below SLO targets.

   Marginal — Storage group performing below SLO target.

   Stable — Storage group performing within the SLO target.
7. Click **OK**.

Editing SLO compliance alerts policies

This procedure explains how to edit SLO compliance alert policies.

Before you begin:

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

To edit compliance alert policies:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › SLO Compliance Alert Policies** to open the **SLO Compliance Alert Policies** list view.
3. Select the policy, and then select (enable) or clear (disable) any of the Compliance States. Unisphere for VMAX only generates alerts for enabled compliance states.
4. Click **Apply**.

Deleting SLO compliance alerts policies

This procedure explains how to delete SLO compliance alert policies.

Before you begin:

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

To delete compliance alert policies:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › SLO Compliance Alert Policies** to open the **SLO Compliance Alert Policies** list view.
3. Select one or more policies and click **Delete**.
4. Click **Delete** in the confirmation message.

Viewing SLO compliance alerts policies

This procedure explain how to view compliance alert polices set on storage systems running HYPERMAX OS 5977 or higher.

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › SLO Compliance Alert Policies** to open the **SLO Compliance Alert Policies** list view.

   The following properties display:

   - **Symmetrix ID** — Storage system on which the policy is defined.
   - **Storage Group** — Storage group on which the policy is defined.
   - **Compliance State** — Enabled compliance states:

     Critical — Storage group performing *well* below SLO targets.

     Marginal — Storage group performing below SLO target.

     Stable — Storage group performing within the SLO target.
   - **Notification** — 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:

   - **Create** — See SLO compliance alert policies on page 80.
   - **Notify** — See Configuring compliance alert notifications on the facing page.
   - **Delete** — See Deleting SLO compliance alerts policies above.
   - **Apply** — Applies policy changes.
   - **Cancel** — Cancels policy changes.

Configuring compliance alert notifications

This procedure explains how to configure Unisphere for VMAX to notify you when a storage group generates a compliance alert.

Before you begin:

- The storage system must be running HYPERMAX OS 5977 or higher.
- The storage system must be configured to deliver alert notifications, as described in Configuring alert notifications on page 48.

To configure compliance alert notifications:

1. From the **System Selector**, select **All Symmetrix**.
2. Select **Home › Administration › Alert Settings › SLO Compliance Alert Policies** to open the **SLO Compliance Alert Policies** list view.
3. Select one or more policies and click **Notify** to open the **Define Notification** dialog box.
4. Select (enable) the method you want to use to deliver the notifications. Possible methods are:

   Note The storage system must already be configured to deliver alerts in the desired method, as described in Configuring alert notifications on page 48.

   - **Enable Email** — Forwards alert notifications to an email address.
   - **Enable SNMP** — Forwards alert notifications to a remote SNMP listener.
   - **Enable Syslog** — Forwards alert notifications to a remote syslog server.
5. Click **OK**.

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.

Note Runtime check alerts will not be sent in alert notifications. For more information on alert notifications, refer to Configuring alert notifications on page 48.

Note Runtime alerts are not storage system-specific and therefore cannot be acknowledged or deleted.

Runtime check	Number of volumes	Threshold	Alert
Total memory on the	0 -	12 GB	System memory *‹# GB›* is below the

Runtime check	Number of volumes	Threshold	Alert
Unisphere server	64,000		minimum requirement of *‹# GB›*
	64,000 - 128,000	16 GB	
	128,000 - 256,000	20 GB	
Free disk space on the Unisphere installed directory	0 - 64,000	100 GB	Free disk space ‹# GB› is below the minimum requirement of *‹# GB›*
	64,000 - 128,000	140 GB	
	128,000 - 256,000	180 GB	
Number of managed storage systems	0 - 64,000	10	Number of managed arrays ‹#› is over the maximum supported number of #
	64,000 - 128,000	20	
	128,000 - 256,000	20	
Number of managed volumes	256,000		Number of managed volumes ‹#› is over the maximum supported number of ‹#›. 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 (*SymmID*) is below the minimum requirement of 6.

Security

Understanding user authorization

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

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

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

Roles

The following lists the available roles. Note that you can assign up to four of these roles per authorization rule. For a more detailed look at the permissions that go along with each role, see Roles and associated permissions on page 98.

- **None** — Provides no permissions.
- **Monitor** — Performs read-only (passive) operations on a storage system.
- **StorageAdmin** — Performs all management (active or control) operations on a Symmetrix system and modifies GNS group definitions in addition to all Monitor operations.
- **Admininstrator** — Performs all operations on a storage 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.
- **DSA Admin** — Collect and analyze database activity with Database Storage Analyzer.

  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 22.

Individual and group roles

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

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.

Viewing user sessions

This procedure explains how to view active user sessions for a storage system.

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › User Sessions** to open the **User Sessions** page.

   The following properties display:

   - **User Name** — Name of the individual or group.
   - **Start Time** — Date and time that the user logged in to the console.
   - **IP Address** — Address of the console.

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 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. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Authentication** to open the **Authentication** page.
3. Click **Edit** to open the **Configure Authentication** wizard.
4. 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.
5. 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.
6. Click **Next**.
7. If you are configuring LDAP-SSL, do the following:
   a. Specify values for the following parameters and click **Next**.
      - **Server (IP or Hostname)** — 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.
      - **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.

b. *Optional*: To upload an SSL certificate, click **Browse**, locate the certificate, and click **Open**. To view the contents of the certificate, click **View Certificate**.

c. *Optional*: To limit authentication to only members of specific LDAP groups, click **Show Advanced**, select the option, and then type the **Group Names**, separated by commas.

d. Click **Next**.

8. 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.
9. Click **Finish**.

Viewing authentication authorities

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Authentication** to open the **Authentication** page.

   Use the **Authentication** page 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.

   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 88.

Authorization

Adding authorization rules

Before you begin:

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

To add authorization rules:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Authorized Users & Groups** to open the **Authorized Users & Groups** list view.
3. Click **Add** to open the **Add Authorization Rule** dialog box.
4. *Optional*: Select an authentication **Authority**. Possible values are:
   - **Local Directory** — Specifies to authenticate the user against the Local Authority repository.
   - **Windows AD** — Specifies to authenticate the user against the Active Directory domain.
   - **LDAP-SSL** — Specifies to authenticate the user against an LDAP directory.
5. Do the following depending on the authority:

   Local Directory:

   Select the user **Name**.

   Windows AD or **LDAP-SSL**

   a. Specify whether the rule is for an individual **User** or for a user **Group**.

   b. *Optional*: Type the name of the **Domain** used to authenticate the user/group. 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

   c. Type the **Name** of the user or group. User/group names can only contain alphanumeric characters.
6. For **Database Storage Analyzer** select **None**, **Read Only**, or **Admin**.By default, **Database Storage Analyzer** permissions are set to **None**.
7. On the **Roles** tab, select the object and up to four roles.
8. Click the **Show only VMAX with hinting enabled** checkbox to see the list of VMAXs with hinting enabled. The VMAX must be running HYPERMAX 5977 or higher.
9. Click the **DSA Hinting** role. This role is only enabled for DSA administrators.
10. Click **OK**.

Modifying authorization rules

Before you begin:

To perform this operation, you must be the Initial Setup User (set during installation), or a SecurityAdmin on all authorized Symmetrix systems.

To modify authorization rules:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Authorized Users & Groups** to open the **Authorized Users & Groups** list view.
3. Select a rule and click **Modify** to open the **Modify Authorization Rule** dialog box.
4. On the **Roles** tab, add or remove roles from any of the available objects, being sure to not exceed the four roles/object limit.
5. Click **OK**.

Removing authorization rules

Before you begin:

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

To remove authorization rules:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Authorized Users & Groups** to open the **Authorized Users & Groups** list view.
3. Select the rule and click **Remove**.
4. Click **OK** on the confirmation message.

Viewing authorization rules

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Authorized Users & Groups** to open the **Authorized Users & Groups** list view.

   Use the **Authorized Users & Groups** list view to view and manage authorization rules.

   The following properties display:

   - **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

   The following controls are available:

   - **Add** — Authorization on page 90
   - **Remove** — Removing authorization rules on the previous page
   - **Modify** — Modifying authorization rules on the previous page

Local Users

Creating local users

This procedure explains how to create local users. Local users have accounts stored locally in the user database the Unisphere server host.

Before you begin:

To perform this operation, you must be the Initial Setup User (set during installation), or SecurityAdmin on at least one Symmetrix system.

To create local users:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Local Users** to open the **Local Users** list view.
3. Click **Create** to open the **Add Authorization Rule** dialog box.
4. Type a **User Name**. User names are case-sensitive and can only contain alphanumeric characters.
5. *Optional*: Type a **Description**.
6. Type and confirm a user **Password**. Passwords cannot exceed 16 characters. Passwords can contain alphanumeric characters but not the following special characters: @ : ? ; " | ‹ › [ ] + = , * \ /
7. For **Database Storage Analyzer** select **None**, **Read Only**, or **Admin**.By default, **Database Storage Analyzer** permissions are set to **None**.
8. Click the **Roles** tab, select the object and up to four roles.
9. Click the **Show only VMAX with hinting enabled** checkbox to see the list of VMAXs with hinting enabled. The VMAX must be running HYPERMAX 5977 or higher.
10. Click the DSA Hinting role. This role is only enabled for DSA administrators.
11. Click **OK**.

Editing local users

This procedure explains how to edit the roles associated with a user or group.

Before you begin:

- To perform this operation, you must be the Initial Setup User (set during installation), or SecurityAdmin on all authorized Symmetrix systems.
- Users cannot remove the SecurityAdmin role from themselves.

To create local users:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Local Users** to open the **Local Users** list view.
3. Select a user and click **Edit** to open the **Properties** dialog box.
4. *Optional*: Type a new **Description**.
5. On the **Roles** tab, add or remove roles from any of the available objects, being sure to not exceed the four roles/object limit.
6. Click **OK**.

Deleting local users

This procedure explains how to delete local users and all fully-qualified authorization rules (rules in the format (L:*HostName**UserName*).

Before you begin:

- To perform this operation, you must be the Initial Setup User (set during installation), or SecurityAdmin on all authorized Symmetrix systems.
- Users cannot remove the SecurityAdmin role from themselves.

To delete a user:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Local Users** to open the **Local Users** list view.
3. Select the user and click **Delete** to open the **Delete Local User** dialog box.
4. *Optional*: By default, deleting local users removes their associated authorization rules from all objects, including storage systems and the Database Storage Analyzer database. To override this default behavior, select the **Remove associated authorization rules** option.
5. 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 Initial Setup User (set during installation), or SecurityAdmin on at least one Symmetrix.

To change local directory user passwords:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Local Users** to open the **Local Users** list view.
3. Select a user and click **View Details** to open the user's **Details** view.
4. Click **Change Password** to open the **Change Password** dialog box.
5. Type the user's **Current Password**.
6. Type a **New Password** and **Confirm New Password**.
7. Click **OK**.

Viewing local users

Before you begin

To perform this operation, you must be the Initial Setup User (set during installation), or Monitor on at least one Symmetrix system.

To view users with a local Unisphere account:

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › 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.

   The following controls are available:

   - **Create** — Local Users on page 93
   - **View Details** —Viewing local users details on the facing page
   - **Delete** — Deleting local users on the previous page
   - **Edit** — Editing local users on page 94

Viewing local users details

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

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Local Users** to open the **Local Users** list view.
3. Select a user and click **View Details** to open its **Details** view.

   The **Details:Local User** 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** — Deleting local users on page 95
- **Change Password** — Changing local user passwords on page 95
- **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.

Viewing local user authorization rules

This procedure explains how to view the authorization rules associated with local users.

1. Select **All Symmetrix**.
2. Select **Home › Administration › Security › Local Users** to open the **Local Users** list view.
3. Select the user and click **View Details** to open its details view.
4. In the **Related Objects** panel, click **Authorization Rules** to open the authorization rules view.

   The following properties display:

   - **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

   The following controls are available:

   - **Add** — Authorization on page 90
   - **Remove** — Removing authorization rules on page 92
   - **Edit** — Modifying authorization rules on page 92

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 storage system until an Administrator or SecurityAdmin is added to the storage system.

Permissions	Roles							
	Administrator	StorageAdmin	Monitor	SecurityAdmin	Auditor	None	Perf Monitor	Database Storage Analyzer Admin
Create/delete user accounts	Yes	No	No	Yes	No	No	No	No
Reset user password	Yes	No	No	Yes	No	No	No	No
Create roles	Yes	Yes	No	Yes (self excluded)	No	No	No	No
Change own password	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Manage storage systems	Yes	Yes	No	No	No	No	No	No
Discover storage systems	Yes	No	No	Yes	No	No	No	No
Add/show license keys	Yes	Yes	No	No	No	No	No	No
Set alerts and Optimizer monitoring options	Yes	Yes	No	No	No	No	No	No
Release storage system locks	Yes	Yes	No	No	No	No	No	No
Set Access Controls	Yes	Yes	No	No	No	No	No	No
Set replication and	Yes	Yes	No	No	No	No	No	No

Permissions	Roles							
	Administrator	StorageAdmin	Monitor	SecurityAdmin	Auditor	None	Perf Monitor	Database Storage Analyzer Admin
reservation preferences								
Access performance data	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Start data traces	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Set performance thresholds/alerts	Yes	Yes	No	No	No	No	Yes	No
Create and manage performance dashboards	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Collect and analyze database activity with Database Storage Analyzer	No	No	No	No	No	No	No	Yes

Link and launch

Creating link-and-launch client registrations

This procedure explains how to register other applications with the SMAS server. Once registered, users of the registered applications can launch 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 Unisphere for VMAX from their respective clients, must have an Administrator or StorageAdmin role assigned to them within Unisphere for VMAX.
- Link-and-launch is not supported with X.509 certificate-based user authentication.

To create link-and-launch client registrations:

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

Editing link-and-launch client registrations

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

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. Select **Home › Administration › Link and Launch** to open the **Link and Launch** list view.
3. Select a registration, and click **Edit** to open the **Edit Launch Client Registration** dialog box.
4. Type the **Current Password**.
5. Type the **New Password**. Passwords can be up to 75 alphanumeric characters.
6. Retype the new password to confirm it.
7. Click **OK**.

Deleting link-and-launch client registrations

This procedure explains how to delete an application registration from the SMAS server. Once deleted, users of the application will no longer be able to launch 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. Select **Home › Administration › Link and Launch** to open the **Link and Launch** list view.
3. Select a registration, and click **Delete**.
4. Click **OK** in the confirmation box.

Viewing link and launch client registrations

1. From the system selector, select **All Symmetrix**.
2. Select **Home › Administration › 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 the previous page
   - **Edit** — Editing link-and-launch client registrations on the previous page
   - **Delete** — Deleting link-and-launch client registrations above

CHAPTER 4
Storage Management

This chapter contains the following:

VMAX provisioning 104
Storage groups 117
Fully Automated Storage Tiering 165
Volume management 208
Disk groups 245
Virtual Provisioning 251
Understanding Virtual LUN Migration 277
External Storage 282
Understanding storage templates 296
Reservations 302
Understanding Optimizer 306
Managing vVols 312

VMAX provisioning

Provisioning VMAX storage (HYPERMAX OS 5977 and higher)

With the release of HYPERMAX OS 5977 and the next generation VMAX Family systems, Unisphere for VMAX introduces support for Service Level Objective provisioning. Service Level Objective provisioning simplifies VMAX management by automating many of the tasks associated with provisioning storage.

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

By default, VMAX systems running HYPERMAX OS 5977 or higher are pre-configured with a single Storage Resource Pool (SRP) containing all the physical disks on the system organized into disk groups by technology, capacity, rotational speed, and RAID protection type. Unisphere for VMAX allows storage administrators to view all the SRPs configured on the system and the demand that storage groups are placing on them. For more information, refer to Managing storage groups (HYPERMAX OS 5977 or higher) on page 117.

In addition, VMAX systems are also pre-configured with a number of Service Level and workloads, which storage administrators use to specify the performance objectives for the application they are provisioning.

When provisioning storage for an application, storage administrators assign the appropriate SRP, Service Level Objective, and workload to the storage group containing the application's LUNs.



Unisphere for VMAX provides the following methods for provisioning 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 (that is, 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 host initiators and mapped to the ports in the port group.

Before you begin:

- The VMAX system must already be configured.
- For instructions on provisioning VMAX storage systems running Enginuity 5874 - 5876, refer to Provisioning VMAX storage (Enginuity 5874 - 5876) on page 108
- For instructions on provisioning DMX storage systems, refer to Masking volumes (page 369).

To provision VMAX storage:

Recommended

1. Create a host

   Use the **Create Host** dialog box 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.

Advanced

1. Create a host

   Use the **Create Host** dialog to group host initiators (HBAs).

2. Create the volumes

   Create one or more volumes on the VMAX system.

3. Create the storage group

   Use the **Create Storage Group** dialog box to add the volumes you just created to a storage group, and associate the storage group with a storage resource pool, a Service Level Objective, and a workload.

4. Create a port group

   Group Fibre Channel and/or iSCSI front-end directors.

5. Create a masking view

   Associate the host, storage group, and port group into a masking view.

Unisphere for VMAX Help: Storage management › VMAX provisioning › Using the Provision Storage wizard (HYPERMAX OS 5977 or higher)

Using the Provision Storage wizard (HYPERMAX OS 5977 or higher)

This procedure explains how to use the Provision Storage wizard to provision storage systems running HYPERMAX OS 5977. In addition, you can also use a subset of the steps to simply create a storage group, without actually provisioning it.

For instructions on provisioning storage systems running Enginuity 5874 - 5876, refer to Using the Provision Storage wizard (Enginuity 5874 - 5876).

Before you begin:

- The user must have Administrator or StorageAdmin permission.
- There are multiple ways to open the **Provision Storage** wizard. Depending on the method you use, some of steps listed below may not apply. For example, if you open the wizard from the **Hosts** view, the step on selecting a host will not apply. Or, if you open the wizard from the Reference Workload view, the steps on selecting the Service Level and Workload Type will not apply.

To use the Provisioning Storage wizard:

1. Select the storage system.
2. Open the Provision Storage to Host wizard by doing either of the following:
   - Hover the cursor over the **Storage** menu and select **Provision Storage to Host** under **Common Tasks** to open the **Provision Storage** wizard.
   - On the **Storage Group Dashboard** (**Storage › Storage Group Dashboard**), open one of the storage group list views (**Total**, **Stable**, **Marginal**, **Critical**, or **No SLO**), and click **Create SG** in the corresponding view.
3. Type a **Storage Group Name** name. 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.
4. Select a **Storage Resource Pool**. To create the storage group outside of FAST control, select **None**. External storage resource pools are listed below the **External** heading.
5. Select the **Service Level** to set on the storage group. Service levels specify the characteristics of the provisioned storage, including average response time, workload type, and priority. This field defaults to None if you set the **Storage Resource Pool** to None. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

6. Select the **Workload Type** to assign it. Workload types are used to refine the service level (that is, narrow the latency range). Possible values are OLTP or DSS, where OLTP workload is focused on optimizing performance for small block I/O and DSS workload is focused on optimizing performance for large block I/O. The **Workload Type** can also specify whether to account for any overhead associated with replication (OLTP_Rep and DSS_Rep, where denotes replication). To provision the storage group using a previously created Reference Workload, select the Reference Workload, in which case, the Service Level will update to reflect parameter defined in the Reference Workload.
7. Type the number of **Volumes** and select the **Capacity** of each.

   Note The maximum volume size supported on a VMAX3 system is 64 TB.

8. *Optional*: Hover the cursor on the service level and click edit ( ) to set any of the following options:
   - To name the volumes you are adding to the storage group, select one of the following **Volume Identifiers** and type a **Name**:

     Note When using this option, Unisphere for VMAX will only use new volumes when creating the storage group; it will not use any existing volumes in the group.

     - **None** — Allows the system to name the volumes (Default).
     - **Name Only** — All volumes will have the same name.
     - **Name and VolumeID** — All volumes will have the same name with a unique volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
     - **Name and 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 **Allocate capacity for each volume** you are adding to the storage group, select the option. You can only use the this option for newly created volumes, not existing volumes.

     If you selected to allocate capacity in the previous step, you can mark the allocation 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.
9. *Optional*: To add a child storage group, click **Add Service Level** and specify a storage group **Name**, **Service Level**, **Workload Type**, **Volumes** and **Capacity**. Repeat this step for each additional child storage group. The maximum number of allowed child storage groups is 64.
10. *Optional*: To set host I/O limits for the storage groups, click **Set Host I/O Limits** to open the **Host I/O Limits** dialog box. For instructions setting the limits, refer to dialog's help page. When done, close the dialog box to return to the wizard.
11. To create a storage group, without actually provisioning it, click one of the following; otherwise, click **Next** and continue with the remaining steps in this procedure:
    - **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 and Previewing.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
12. Specify the host/host group to use by selecting an existing host/host group, or doing the following to create a new host/host group. When done, click **Next**.
    - To create a new host, click **Create Host** to open the **Create Host** dialog box. For instructions on creating a host, refer to the dialog's help page.

- To create a new host group, click **Create Host Group** to open the **Create Host Group** dialog box. For instructions on creating a host, refer to the dialog's help page.

13. Select whether to use a **New** or an **Existing** port group, and then do the following depending on your selection:

    New:

    a. *Optional*: Edit the suggested **Port Group Name** by highlighting it and typing a new name over it. Port group names must be unique from other port groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Port group names are case-insensitive.

    b. Select the ports to use. To view host-invisible ports (unmasked and unmapped), select **Include ports not visible to the host**. If a Fibre or iSCSI host was not selected, select the appropriate filter to filter the port list by iSCSI virtual ports or FC ports based on the selected host. If an empty host was selected, the radio button Fibre is selected by default. The Dir-Port table is filtered to only show either FC or iSCSI depending on the radio button selection.

    The following properties display:

    - **Dir-Port** — Storage system director and port in the port group.
    - **Identifier** — Port identifier.
    - **Initiators Logged In** — Number of initiators logged into the fabric.
    - **PGs** — Number of port groups where the port is a member.
    - **Mappings** — Number of mappings.
    - **% Busy** — Percentage of time that the port is busy.

    c. Click **Next**.

    Existing: Select the port group and click **Next**.

14. *Optional*: Edit the suggested name for the **Masking View** by highlighting it and typing a new name over it. Verify the rest of your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
15. *Optional*: To receive alerts when the performance of the storage group changes, relative to its SLO target, select **Enable SLO Compliance Alerts**. For more information on SLO Compliance Alerts, refer to Creating SLO compliance alerts policies.
16. *Optional*: Determine if the storage system can handle the updated service level:

    a. Click **Run Suitability Check**. The **Suitability Check** dialog box opens, indicating the suitability of the change. For information on interpreting the results, refer to the dialog's help page. This option is only available under certain circumstances. For more information, refer to Suitability Check restrictions.

    b. Click **OK** to close the message.

    c. If your updates are found to be unsuitable, modify the settings and run the check again until the suitability check passes.

17. Do one of the following:
    - 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 and Previewing.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Provisioning VMAX storage (Enginuity 5874 - 5876)

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

Unisphere for VMAX provides the following methods for provisioning 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 (that is, 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 host initiators and mapped to the ports in the port group.



Before you begin:

- The VMAX system must already be configured.
- For instructions on provisioning DMX storage systems, refer to Masking volumes (page 369).

To provision VMAX storage:

Recommended

1. Create a host

   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.

Advanced

1. Create a host

   Use the Create Host wizard to group host initiators (HBAs)
2. Create the volumes

   Create one or more volumes on the VMAX 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 storage groups.

4. Create a port group

   Group Fibre Channel and/or iSCSI front-end directors.

5. Create a masking view

   Associate the host, storage group, and port group 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 Provision Storage wizard (Enginuity 5874 - 5876)

This procedure explains how use the Provision Storage wizard to provision storage systems running Enginuity 5874 - 5876. 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 (that is, the ability to create their own volumes, storage groups, etc.). In addition, you can also use a subset of the steps to simply create a storage group, without actually provisioning it.

For users who want the ability to control every aspect of the provisioning process, refer to the **Advanced** procedure in Provisioning VMAX storage (Enginuity 5874 - 5876).

For instructions on provisioning storage systems running HYPERMAX OS 5977, refer to Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106.

Before you begin:

- The storage system must already be configured and you must already have a host. For instructions on creating a host, refer to See Creating hosts on page 322..
- For instructions on provisioning storage systems running Enginuity 5671 - 5774, refer to Masking volumes (page 369).
- Note the following recommendations:
  - Port groups should contain four or more ports.
  - Each port in a port group should be on a different director.
- There are multiple ways to open the **Provisioning Storage** wizard. Depending on the method you use, some of the following steps may not apply. For example, selecting a storage group in the Storage Groups list view and clicking Provision Storage to Host will open the wizard on the Select Host/Host Group page because you are starting out with a storage group.

To use the Provisioning Storage wizard:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Provision Storage to Host** under **Common Tasks** to open the **Provision Storage** wizard.
3. Type a **Storage Group Name**. 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
4. Select the **Storage Group Type**.
5. Do the following, depending on the storage group type:
   Standard Storage Group:
   a. Select the **Volume Type** to add to the storage group and click **Next**.
   b. Do the following, depending on the volume type:
      Virtual Volumes:
      1. Select the **Emulation** type for the volumes to add to the storage group.
      2. *Optional*: Select the **Thin Pool** containing the volumes to add to the storage group.
      3. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**.
      4. *Optional*: Hover the cursor over the volumes and click edit to set any number of the following options. Once done, click **OK** and continue with the remaining steps in this procedure.
         - 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.

           - **None** — Allows the system to name the volumes (Default).
           - **Name Only** — All volumes will have the same name.
           - **Name and VolumeID** — All volumes will have the same name with a unique storage volume ID appended to them. When using this option, the maximum number of characters allowed is 50.
           - **Name and 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 set an **SSID** for the volumes, type a value or click **Select** to open a dialog box from which you can select an SSID.
         - To only use BCVs in the storage group, select **Use BCV volumes**.

- To **Allocate capacity for each volume** you are adding to the storage group, select the option.

  If you selected to allocate capacity in the previous step, you can mark the allocation 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.

Regular Volumes:

1. Select the **Disk Technology** on which the storage group will reside.
2. Select the **Emulation** type for the volumes to add to the storage group.
3. Select the **Protection** level for the volumes to add to the storage group.
4. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Size**. To create an empty storage group, leave both of these fields empty.
5. *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 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.

     - **None** — Allows the system to name the volumes (Default).
     - **Name Only** — All volumes will have the same name.
     - **Name + VolumeID** — All volumes will have the same name with a unique storage 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.

Template:

1. Select the name of the template.
2. Type the **Number of Volumes** to add to the storage group.
3. *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

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 storage 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.

Cascaded Storage Group:

a. Click **Next**.

b. Select one or more child storage groups.

Empty Storage Group:

Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

1. If you want to create a storage group, without actually provisioning it, click one of the following; otherwise, click **Next** and continue with the remaining steps in this procedure:
   - **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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
2. Specify the host/host group to use by selecting an existing host/host group, or doing the following to create a new host/host group.
   - To create a new host, click **Create Host** to open the **Create Host** dialog box. For instructions on creating a host, refer to the dialog's help page.
   - To create a new host group, click **Create Host Group** to open the **Create Host Group** dialog box. For instructions on creating a host, refer to the dialog's help page.
3. Click **Next**.
4. Select whether to use a **New** or an **Existing** port group, and then do the following depending on your selection. When done, click **Next**.

   New:

a. *Optional*: Edit the suggested port group **Name** by highlighting it and typing a new name over it. Port group names must be unique from other port groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Port group names are case-insensitive.

b. Select the ports to use. To view host-invisible ports (unmasked and unmapped), select **Include ports not visible to the host**.

   The following properties display:

   - **Dir-Port** — Storage system director and port in the port group.
   - **Initiators Logged In** — Number of initiators logged into the fabric.
   - **# Port Groups** — Number of port groups where the port is a member.
   - **# Mappings** — Number of volumes mapped to the port.

c. Click **Next**.

Existing: Select the port group and click **Next**.

5. *Optional*: Edit the suggested name for the **Masking View** by highlighting it and typing a new name over it. Verify the rest of your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Reference workloads

Creating Reference Workloads

This procedure explains how to create Reference Workloads. Reference Workloads are a reusable set of workload parameters that simplify provisioning by eliminating the ambiguity associated with matching a new workload with the appropriate pre-defined workload. Using this feature, Storage administrators create Reference Workloads from known pre-configured workloads, and then associate them with new workloads during the provisioning process.

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- This feature is only available for a user with Admin or StorageAdmin permissions.

To create Reference Workloads:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **SLO Compliance** panel, click **Total** to open the **Storage Groups** list view.

4. Select the storage group containing the workload that you want to reference and click **View details** to open its details view.
5. If not already displaying, click the **SLO Compliance** tab.
6. Click **Save as Reference Workload** to open the **Create Reference Workload** dialog box.
7. Type a **Reference Workload Name**. Reference Workload names must be unique and cannot exceed 64 characters.
8. Click **OK**.

Renaming Reference Workloads

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- This feature is only available for a user with Admin or StorageAdmin permissions.

To rename Reference Workloads:

1. Select the storage system.
2. Select **Storage › Service Levels** to open the **Service Levels** view.
3. Click the **Workload Types** tab.
4. Hover the cursor over the Reference Workload and click Edit ( ).
5. Type a new name over the existing. Reference Workload names must be unique and cannot exceed 64 characters.
6. Click **OK**.

Viewing Reference Workloads

This procedure explains how to view Reference Workloads.

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view Reference Workloads:

1. Select the storage system.
2. Select **Storage › Service Levels** to open the **Service Levels** view.
3. Click the **Workload Types** tab.

   The following properties display:

   - **Reference Workloads** — Workloads configured on the system, including:

     Standard — Pre-configured workloads.

     Custom — User-configured workloads.
   - **Details** — Displays the following details for the selected workload:

- **Service Level** — Service level set for the workload. Service levels specify the characteristics of the provisioned storage, including average response time, workload type, and priority. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

- **Workload** — Workload type set for the workload. Workload types are used to refine the service level (that is, narrow the latency range) Possible values are OLTP or DSS, where OLTP workload is focused on optimizing performance for small block I/O and DSS workload is focused on optimizing performance for large block I/O. The Workload Type can also specify whether to account for any overhead associated with replication (OLTP_Rep and DSS_Rep, where denotes replication).
- **I/O Density** — How efficiently FAST is managing I/O for the workload.
- **Skew** — Calculated skew density score for the disks in the storage group as a percentage of the storage group's expected values.
- **I/O Mixture** — I/O mixture for the workload.

The following control is available:

Provision Storage to Host — Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106. When opening the Provision Storage Wizard from this view, the steps on selecting the Service Level and Workload Type will not apply, as they will be pre-populated with information from the Reference Workload.

Storage groups

Managing storage groups (HYPERMAX OS 5977 or higher)

The **Storage Groups Dashboard** provides you with a single place to monitor and manage storage groups.

Before you begin:

- The storage system must be running HYPERMAX OS 5977 or higher.
- To refresh the information displayed in the **Storage Groups Dashboard**, click refresh Symmetrix ( ) in the status bar, or exit/return to the **Storage Groups Dashboard**.

To access the Storage Groups Dashboard:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** .

Storage Groups Dashboard

The **Storage Groups Dashboard** is organized into the following panels:

SLO Compliance panel

Displays how well the system's storage groups are complying with their respective service level policies, if applicable. All of the storage groups on the VMAX system are organized into the following categories:

Total — All storage groups on the storage system.

29 Stable — Number of storage groups performing within the SLO targets. indicates that there are no storage groups performing within the SLO targets.

10 Marginal — Number of storage groups performing below SLO targets. indicates that there are no storage groups performing below SLO targets.

1 Critical — Number of storage groups performing *well* below SLO targets. indicates that there are no storage group performing well below SLO targets.

9 No SLO — No SLO compliance information.

Storage Group Management panel

Displays the following links:

- **Provision Storage to Host** — Opens the Provision Storage to Host wizard, which will guide you through the process of provision storage to a host.
- **View SLO Compliance Report** — Opens the SLO Compliance report, which displays storage group performance against Service Level Objectives over a period of time. This option is only available when performance is enabled and Workload Planner has processed.

In addition, this panel also includes a **Search** widget, which allows you to quickly locate a specific storage group on the storage system. Double-clicking the resulting storage group will open its **Details** view.

Storage Resource Pools panel

Displays performance information for the Storage Resource Pools (SRPs) on the storage system. By default, VMAX systems only include a single SRP, the DEFAULT_SRP. In the event that there are multiple SRPs, each will have their own tab in the display.

The following information displays for each pool:

- Alert summary icon that displays the status and number of alerts for the SRP. Alert status is indicated by icon color:

  One or more fatal alerts

  One or more critical alerts, with none higher

  One or more warning alerts, with none higher

  One or more informational alerts, with none higher

  No alerts

  To view details on the alerts, click the icon to open the SRP Alert list view.

- **Capacity** — Displays used (allocated) verses free capacity for the SRP, with the option to overlay the amount of subscribed capacity (**Display subscription**). You also have the option to view only **SRDF DSE allocated** capacity for the SRP or **Snapshot allocated** capacity for the SRP.

- **Headroom** — Displays the space available for a particular SLO/workload combination if all remaining capacity was on that type.
  For external SRPs, headroom information is not displayed.

- **Demand Report** — Displays the allocated and subscribed capacity for the SRP by service level:
  - **SLO** — Name of the Service Level Objective.
  - **Subscription** — Total space the SLO has subscribed on the SRP in GB.
  - **%** — Percentage of space the SLO has subscribed on the SRP.
  - **Allocated (GB)** — Total space the SLO has used on the SRP in GB.

- **Reports** — Displays links to the following reports:
  - **Storage Group Demand Report** —Displays all the storage groups on the SRP and their associated demand.
  - **Workloads Demand Report** — Displays demand that each workload is placing on the SRP.
  - **View SRP Details** — Opens the **SRP Details** view, from which you can view and manage the SRP. For more information, refer to Viewing Storage Resource Pools on page 168.

Creating storage groups (HYPERMAX OS 5977 or higher)

This procedure explains how to create storage groups on storage systems running HYPERMAX OS 5977 or higher. In addition to method described below, you can also create a storage group using the Provision Storage wizard, as described in Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106

For instructions on creating storage groups on storage systems running Enginuity 5874 - 5876, refer to Creating storage groups (Enginuity 5874 - 5876) on the next page.

Before you begin:

- The user must have Administrator or StorageAdmin permission.
- The maximum number of storage groups allowed on a storage system running HYPERMAX OS 5977 is 16,384.
- A storage group can contain up to 4,096 volumes.
- A volume can belong to multiple storage groups if only one of the groups is under FAST control.

To create a storage group:

1. Select the storage system.
2. Select **Storage › Volumes** to open the **Volumes** list view.
3. Select one or more volumes and click **Create SG** to open the **Create Storage Group** dialog box.
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 (-) are allowed. Storage group names are case-insensitive.
5. To create the storage group outside of FAST control, set **Storage Resource Pool** to **None**; otherwise, leave this field set to the default.
6. Select the Service Level Objective (**SLO**) to set on the SG. Service level policies specify the characteristics of the provisioned storage, including maximum response time, workload type, and priority. This field defaults to None if you set the **Storage Resource Pool** to None. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

7. Refine the SLO by selecting the **Workload** to assign it.
8. Click **OK** to create the storage group now, or click **Show Advanced** to continue setting the advanced options, as describe in the remaining steps.
9. *Optional*: Specify to **Allocate Full Volume capacity**.
10. If you selected to allocate capacity in the previous step, you can mark the allocation 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.
11. Click **OK**.

Creating storage groups (Enginuity 5874 - 5876)

This procedure explains how to create storage groups on VMAX systems running Enginuity 5874 - 5876. For instructions on creating storage groups on VMAX systems running HYPERMAX OS 5977, refer to Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106.

Before you begin:

- You 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 231
- This procedure explains how to create storage groups using the **Provision Storage** wizard, which involves only the first page of the wizard. For instructions on completing the wizard, thereby provisioning the new storage groups, refer to Using the Provision Storage wizard (Enginuity 5874 - 5876) on page 110.

To create a storage group:

1. Select the storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Click **Create SG** to open the **Provision Storage** wizard.
4. Refer to wizard help system for instructions.

Adding volumes to storage groups

This procedure explains how to add volumes to existing storage groups.

Before you begin:

Storage groups require Enginuity 5874 - 5876, or HYPERMAX OS 5977 or higher.

To add volumes to storage groups:

1. Select the storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **SLO Compliance** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and click **View Details** to open the its **Details** view.
2. In the **Related Objects** panel, click **Volumes** to open the storage group **Volumes** list view.
3. Click **Add Volumes to SG** to open the **Add Volumes to Storage Group** wizard.
4. Locate the volumes by selecting/typing values for any number of the following criteria:

- **Capacity equal to** — Filters the list for volumes with a specific capacity.
- **Volume ID** — Filters the list for a volume with specific ID.
- **Volume Identifier Name** — Filters the list for the specified volume name.
- **Additional Criteria** — Allows you to refine your query by adding conditional expressions, as describe next:
  a. Select values for the following:
     - **Category** — Specifies the broad category.
     - **Attribute** — Refines the category.
     - **Operator** — Establishes the relationship between the attribute and the value.
     - **Value** — Conditions for the specified attribute.
  b. To add another expression, click **Add Another**, and repeat the previous step. To clear your criteria, click **Clear All**.
- **Include Volumes in Storage Groups** — Includes volumes already in other storage groups.

5. Click **Find Volumes** to run the query. Results are displayed on the next page in the wizard.
6. Do one of the following:
   - Select the **Volumes** and click **Add to SG**.
   - Click **Modify Criteria** to modify your search query.

Copying volumes between storage groups

This procedure explains how to copy volumes between storage groups.

Before you begin:

- Storage groups require Enginuity 5874 - 5876, or HYPERMAX OS 5977 or higher.
- The user must have StorageAdmin permission.

To copy volumes between storage groups:

1. Select the storage system.
2. Do the following depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and click **View Details** to open its **Details** view.
2. In the **Related Objects** panel, click **Volumes** to open the storage group **Volumes** list view.
3. Select one or more volumes click more >> and click **Copy Volumes To SG** to open the **Copy Volumes to Storage Group** dialog box.
4. Select the **Target Storage Group Name**.
5. Click **OK**.

Moving volumes between storage groups

This procedure explains how to move volumes between storage groups.

Before you begin:

- Storage groups require Enginuity 5874 - 5876, or HYPERMAX OS 5977 or higher.
- The user must have StorageAdmin permission.
- To perform this operation without disrupting the host's ability to view the volumes, at least one of the following conditions must be met:
  - Each storage group must be a child of the same parent storage group, and the parent storage group must be associated with a masking view.
  - Each storage group must be associated with a masking view, and both masking views must contain a common initiator group and a common port group. In this scenario, the port groups can be different, but they must both contain the same set of ports, or the target port group can contain a superset of the ports in the source port group.
  - The source storage group is not in a masking view.

To move volumes between storage groups:

1. Select the storage system.
2. Do the following depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and click **View Details** to open its **Details** view.
2. In the **Related Objects** panel, click **Volumes** to open the storage group **Volumes** list view.
3. Select one or more volumes, click more >> , and click **Move Volumes to SG** to open the **Move Volumes to Storage Group** dialog box.
4. Select the **Target Storage Group Name**.
5. *Optional*: By default, the operation will fail if at least one of the conditions above is not met. To override this default behavior, select **Use force flag**.
6. Click **OK**.

Removing volumes from storage groups

This procedure explains how to remove volumes from storage groups.

Before you begin:

Storage groups require Enginuity 5874 - 5876, or HYPERMAX OS 5977 or higher.

To remove volumes from storage groups:

1. Select the storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and click **View Details** to open the its **Details** view.
2. In the **Related Objects** panel, click **Volumes** to open the storage group **Volumes** list view.
3. Select one ore more volumes and click **Remove Volumes** to open the **Remove Volume** dialog box.
4. *Optional*: To unbind the volumes, select **Unbind** or **Unmap**, depending on the storage operating environment.

5. 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Expanding storage groups (Enginuity 5874 - 5876)

This procedure explains how to increase the amount of storage in a group accessible to the masking view or in the FAST Policy.

Before you begin:

- This procedure requires Enginuity 5874 - 5876.
- 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 231

To expand a storage group:

1. Select the storage system.
2. Select **Storage › Storage Groups** to open the **Storage Group** 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**.
- Specify to **Allocate Full Volume capacity**.
- If you selected to allocate full volume capacity in the previous step, you can mark the allocation 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.
- 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.

  - **None** — Allows the system to name the volumes (Default).
  - **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 **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 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.

- **None** — Allows the system to name the volumes (Default).
- **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. Select the **Template** and click **Next**.

b. Type the **Number of Volumes** to add to the storage group.

c. Select the **Thin Pool** containing the volumes to add o the storage group.

d. *Optional*: Click **Show Advanced** and set any of the following options:

- z/OS Only: If the storage group contains CKD volumes, type an **SSID** for the target, or click **Select ...** to open a dialog from which you can select an SSID.
- 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.

- **None** — Allows the system to name the volumes (Default).
- **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**.

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.

- **None** — Allows the system to name the volumes (Default).
- **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. Click **Finish**.

Modifying storage groups

This procedure explains how to modify storage groups on storage systems running HYPERMAX OS 5977 or higher.

This procedure explains how to perform the operation from the **Storage Groups Dashboard**. In eNAS operating environments, you can also perform this operation from **Storage Groups for File** page (**System › Dashboard › File Dashboard › File Storage Groups**).

Before you begin:

- You must be an Administrator or StorageAdmin.
- The maximum number of storage groups allowed on a storage system is 8,192.
- A storage group can contain up to 4,096 storage volumes.
- A volume can belong to more than one storage group.
- A volume can belong to multiple storage groups if only one of the groups is under FAST control.

To modify a storage group:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group and click **Modify** to open the **Modify Storage Group** dialog box.
5. Do any number of the following:
   a. Change the **Storage Group Name** by highlighting it and typing 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.

      Note the following about renaming storage groups:

      - If renaming a storage group with workload on it, you will have to wait some time before the workload is visible in the storage group's **Details** view.
      - When renaming a storage group configured compliance alerts, the compliance alerts will need to be deleted manually. For instructions, refer to Deleting SLO compliance alerts policies on page 82.

   b. Change the **Storage Resource Pool** by selecting the new pool from the drop-down menu. Setting this property to **None** creates the storage group outside of FAST control. External storage resource pools are listed below the **External** heading.

   c. Change the **Service Level** for the storage group. Service level objectives specify the characteristics of the provisioned storage, including maximum response time, workload type, and priority. This field defaults to None if you set the **Storage Resource Pool** to None. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

   d. Change the **Workload Type** assigned to the service level.

   e. Add or remove **Volumes**.

   f. Do the following to change the capacity of the storage group, depending on whether the group contains volumes of the same capacity or mixed capacities:

      - If the group contains volumes of the same capacity, do one of the following:
        - Type or select an increased number of volumes in the **Volumes** drop-down menu.

- Type or select an increased unit capacity of the volumes and/or change the unit in the **Volume Capacity** drop-down menus.

◆ If the group contains volumes of mixed capacities, click **Edit custom capacity** to open the **Modify Custom Capacity** dialog box. Change the number of **Volumes** by capacity, and click **OK**.You can only use the **Allocate capacity for each volume** option for newly created volumes, not existing volumes.

The **Total Capacity** and **Additional Capacity** figures are updated to reflect any changes.

Note The maximum volume size supported on a VMAX3 system is 64 TB.

g. *Optional*: Click **Add Service Level** to add a child storage group. Modify any of the service level parameters, as described earlier in this procedure.

6. *Optional*: To determine if the storage system can handle the updated service level:
   a. Click **Run Suitability Check**. The **Suitability Check** dialog box opens, indicating the suitability of the change. For information on interpreting the results, refer to the dialog's help page. This option is only available under certain circumstances. For more information, refer to Suitability Check restrictions.
   b. Click **OK** to close the message.
   c. If your updates are found to be unsuitable, modify the settings and run the check again until the suitability check passes.
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 851 and Previewing on page 850.
   ◆ Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Moving storage groups to EMC CloudArray

This procedure explains how to move storage groups to EMC CloudArray.

If a child storage group is moved to CloudArray, all sibling storage groups are also moved. If a parent storage group is moved to CloudArray, the child storage groups are also moved.

Before you begin:

◆ The user must have Admin or StorageAdmin permission.

◆ The storage system must have an EMC CloudArray SRP.

To move storage groups to EMC CloudArray:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. Select a storage group, click more >> , and click **Move to EMC CloudArray** to open the **Move to EMC CloudArray** dialog box.

   Information about the storage group and the SRPs from which and to which it is being moved displays.
4. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Protecting storage groups

The **Protect Storage Group** wizard guides you through the process of protecting your storage group. Depending on the capabilities of the storage system, the following options may be available:

- **Snap/VX** — For instructions, refer to Creating snapshots on page 409. This is the default method for storage systems running HYPERMAX OS 5977 or higher.
- **TimeFinder/Clone** — For instructions, refer to Protecting storage groups using TimerFinder/Clone below. This is the default method for storage systems running Enginuity 5671 - 5876.
- **ProtectPoint** — For instructions, refer to Protecting storage groups using ProtectPoint on page 285. This method is only available on storage systems running HYPERMAX OS 5977 or higher.
- **RecoverPoint** — For instructions, refer to Protecting storage groups using RecoverPoint on page 522. This is method is only available for storage systems running Enginuity 5875 - 5876.
- **SRDF** — For instructions, refer to Protecting storage groups using SRDF on page 136. This method is available for storage systems, subject to connectivity rules.
- **SRDF/Metro** — For instructions, refer to Protecting storage groups using SRDF/Metro on page 139. This is method is only available for storage systems running HYPERMAX 5977 and higher..

Protecting storage groups using TimerFinder/Clone

This procedure explains how to clone storage groups.

Before you begin:

- This feature requires the Enginuity 5876.163.105 or higher. This feature does not apply to storage systems running HYPERMAX OS 5977 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_GRPS option must be enabled. For instructions on enabling the option, refer to *Editing the Options file* in the *Solutions Enabler Installation Guide*.
- Meta volumes are not supported.

To protect storage groups using TimeFinder/Clone:

1. Select the storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.
4. If not already selected, select **Point In Time Protection Using Clone**.
5. Click **Next**.
6. Type the name of the device group that will hold the target volumes **(Device Group Name)**.
7. 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.
8. *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 (**Storage Group Name**). Changing the name will also change the target volume storage group name.
9. *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
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 851 and Previewing on page 850.
    - 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 the mode in which to create the clone session **Clone Copy Mode**. You can set the default clone copy mode in the **Preferences** dialog box, as described in Setting system preferences on page 38. The mode you specify here will override the default mode specified in the preferences. Possible values are:
    - **No Copy No Diff** — Create a nondifferential (full) copy session without a full background copy.
    - **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.
   - **Create New** — Specifies to create new volumes to use as targets.
3. Select the type of volumes to use as the targets (**Clone Targets**).
4. Click **Next**.
5. Verify your selections, and then 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Protecting storage groups using ProtectPoint

This procedure explains how to backup thin volumes to Data Domain arrays.

Before you begin:

- The storage system must be running HYPERMAX OS 5977.
- You must have StorageAdmin permission.
- The Data Domain appliance must be connected and zoned to the storage system.
- Provide the Data Domain Admin the number and size of volumes in the production storage group and request that they provide you with double the number of similar volumes (masked/visible to the storage system). For example, if the production storage group contains 10 volumes, the Data Domain Admin should provide you with the LUN numbers of 20 similar volumes.

To protect storage groups using ProtectPoint:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.
5. Select **Backup Using ProtectPoint**.
6. Click **Next**.
7. Type the name of the **Point In Time Copy Name** and click **Next**.
8. Type a name for the **Backup Storage Group**, or leave the system-generated suggestion.
9. Select the external LUNs to add to the backup storage group and click **Add to Storage Group**. Note that the external LUNs you select must match in number and capacity the volumes in the production storage group.
10. Click **Next**.
11. Type a name for the **New Restore Storage Group**, or leave the system-generated suggestion.
12. Select the external LUNs to add to the restore storage group and click **Add to Storage Group**. Note that the external LUNs you select must match in number and capacity the volumes in the production storage group.
13. Click **Next**.
14. Verify your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
15. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
16. Once the job completes successfully, provide the following information to the Data Domain Admin:
    - The LUN numbers used in the backup storage group
    - The LUN numbers used in the restore storage group
    - The name of the point in time copy

Protecting storage groups using SRDF

This procedure explains how to protect storage groups using SRDF.

Before you begin:

- SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977.
- You must have StorageAdmin permission.
- Connectivity to remote storage system must be available.
- All storage systems involved must be discoverable and manageable from the Unisphere for VMAX console.
- The SRDF wizard in Unisphere for VMAX 8.1 and higher releases supports the mandatory creation of a storage group and the optional creation of a device group. The storage group may contain non-concurrent SRDF devices of any one SRDF type, or may contain non-SRDF devices.
- The following validation check is performed by the wizard to determine if selected storage group be SRDF protected: Volumes in the storage group need to be all TDEV's, or all volumes in the storage group need to be R1s and in the same SRDF Group, or all volumes need to be R2s and in the same SRDF Group.

To protect storage groups using SRDF:

1. Select the storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard. The **Select Protection Type** page displays.
2. Select **Remote Replication Using SRDF**.
3. Click **Next**.
4. Select the **Remote Symmetrix** system. To update the list of remote systems, click **Scan**.
5. Select the **Replication Mode**. For more information, refer to SRDF session modes on page 483.
6. Select **Auto** to automatically select a SRDF group or **Manual** to select a SRDF group from a list.
7. *Optional*: To not start pair mirroring, clear the **Establish Pairs** option.
8. Do the following, depending on the storage operating environment (target system):
   For HYPERMAX OS 5977 or higher:
   Optional: Change the **Remote Storage Group Name**, and optionally select a **Remote SLO**. Changing the name will also change the target volume storage group name.
   For Enginuity 5874 - 5876:
   a. Select the **Remote Thin Pool** to which the target volumes will be bound. 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.
   b. *Optional*: Select the **Remote FAST Policy**. This is the FAST policy associated with the remote storage group.
   c. *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.
1. Optional: Click **Device Group** check box and select the **Device Group Name** that will hold the target volumes.
2. Click **Next**. The **Finish** page displays.
3. Verify your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
4. Do one of the following:
   - 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 851 and Previewing on page 850.

- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Protecting storage groups using SRDF/Metro

This procedure explains how to protect storage groups using SRDF/Metro, in order to improve support for host applications in high availability environments.

Before you begin:

- SRDF requires HYPERMAX OS 5977 or higher.
- You must have StorageAdmin permission.
- Connectivity to remote storage system must be available.
- All storage systems involved must be discoverable and manageable from the Unisphere for VMAX console.

To protect storage groups using SRDF Metro:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard. The **Select Protection Type** page displays.
5. Select **High Availability using SRDF/Metro** .
6. Click **Next**.
7. Select the **Remote Symmetrix** system. To update the list of remote systems, click **Scan**.
8. Select **Auto** to automatically select a SRDF group or **Manual** to select a SRDF group from a list.
9. *Optional*: To stop the initiation of pair mirroring, clear the **Establish Pairs** option.
10. If **Establish Pairs** is checked, choose Managed by **Witness** or **Bias**. If Witness is unavailable on the local or remote Symmetrix, the option is disabled and Bias is selected by default. If available, Witness is selected by default.
11. *Optional*: Change the **Remote Storage Group Name**, and optionally select a **Remote SLO**. Changing the name will also change the target volume storage group name
12. Click **Next**. The **Finish** page displays.
13. Verify your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
14. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Converting storage groups to cascaded

This procedure explains how to non-disruptively convert a standalone storage group to cascaded storage group. Once complete, the original storage group will serve as the parent to a new child storage group.

This procedure explains how to perform the operation from the **Storage Groups Dashboard**. In eNAS operating environments, you can also perform this operation from **Storage Groups for File** page (**System › Dashboard › File Dashboard › File Storage Groups**).

Before you begin:

- You must have Administrator or StorageAdmin permission.
- The storage system must be running HYPERMAX OS 5977 or higher.

To convert storage groups:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group, click >>, and select **Convert to Cascaded** to open the **Convert to Cascaded** dialog box.
5. Type a new name over the system-suggested **Child 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 (-) are allowed. Storage group names are case-insensitive.
6. Click **OK**.

Changing Storage Resource Pools for storage groups

This procedure explains how to change the Storage Resource Pool of a parent storage group, with child SLOs using different Storage Resource Pools.

This procedure explains how to perform this operation from the **Storage Groups Dashboard**. In eNAS environments, you can also perform this operation from the **File Storage Groups** page (**System › Dashboard › File Dashboard › File Storage Groups**).

Before you begin:

- The storage system must be running HYPERMAX OS 5977 or higher.
- You must have Administrator or StorageAdmin permission.

To change the Storage Resource Pool for storage groups:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group, click more >> , and select **Change SRP** to open the **Change SRP** dialog box.
5. Select the new **SRP**.
6. *Optional*: Change the **Service Level** for the SG. Service level objectives specify the characteristics of the provisioned storage, including maximum response time, workload type, and priority. This field defaults to None if you set the **Storage Resource Pool** to None. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

7. *Optional*: Refine the SLO by selecting the **Workload Type**to assign it.
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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Associating 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.

Before you begin:

- This feature requires Enginuity 5874 - 5876.
- Storage groups and FAST policies can only be associated under the following conditions:
  - The VMAX system is running Enginuity 5874 through 5876.

- 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 storage 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** panel, click **Manage Policies** to open the **Manage Policies** view.
5. Select the policy and click **Associate Storage Groups** to open the **Associate Storage Groups** dialog box.
6. Select one or more storage groups, and then 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 storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group, click more >> and select **Associate to FAST** to open the **Associate to FAST Policy** dialog box.
4. Select a policy and click **OK**.

Disassociating FAST policies and storage groups

1. Select the storage 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 - nn** 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group, click more >>, and click **Reassociate to FAST Policy** to open the **Reassociate to FAST Policy** dialog box.
4. Select a policy and click **OK**.

Adding or removing cascaded storage groups

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

Before you begin:

Storage groups require Enginuity 5874 - 5876, or HYPERMAX OS 5977 or higher.

To add/remove cascaded storage groups:

1. Select the storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the parent storage group and click **View Details** to open its **Details** view.
2. In the **Related Object** panel, click **Child Storage Group** to open the **Child Storage Groups** list view.
3. Do the following, depending on whether you are adding or removing storage groups:
   Adding storage groups:
   a. Click **Add Storage Group** to open the **Add Storage Group** 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. *Enginuity 5874 - 5976 only*: To unmap, select **Unmap Device**.
   c. Click **OK** in the confirmation message.

Deleting storage groups

This procedure explains how to delete storage groups.

Before you begin:

- Storage groups require Enginuity 5874 - 5876, or HYPERMAX OS 5977 or higher.
- You cannot delete a storage group that is part of a masking view or associated with a FAST Policy.
- Before you can delete a child storage group, you must first remove it from its parent.
- When a storage group configured compliance alerts (requires HYPERMAX OS 5977 or higher) is deleted or renamed, the compliance alerts will need to be deleted manually. For instructions, refer to Deleting SLO compliance alerts policies on page 82.

To delete a storage group:

1. Select the storage system.
2. Do the following depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.

   Note In eNAS operating environments, you can also perform this operation from **Storage Groups for File** page (**System › Dashboard › File Dashboard › File Storage Groups**).

   For Enginuity 5874 - 5876:

   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group, click more >> , and select **Delete**.
2. Click **OK** in the confirmation message.

Setting host I/O limits

Host I/O limits (quotas) 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 storage 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 storage 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 the parent and/or the child storage group. If set for both, the child limits cannot exceed that of the parent.

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.

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 storage system must be running Enginuity 5876.159.102 or higher, or HYPERMAX OS 5977 or higher.
- For Enginuity 5876.159.102 up to HYPERMAX OS 5977, the maximum number of quotas per array is 2,000. For HYPERMAX OS 5977 and higher, the maximum number is 16,000.
- For more information on setting host I/O limits, refer to the *Solutions Enabler Array Management CLI Product Guide*. This guide is part of the *Solutions Enabler Complete Documentation Set*.

To set host I/O limits:

1. Select a storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.

   Note In eNAS operating environments, you can also perform this operation from **Storage Groups for File** page (**System › Dashboard › File Dashboard › File Storage Groups**).

   For Enginuity 5874 - 5876:

   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and select **Set Host I/O Limits** to open the **Set Host I/O Limits** dialog box.
2. 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.
3. 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 Enginuity 5876.163.105 or higher.
   - **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.

- **OnFailure** — 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.

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 851 and Previewing on page 850.
   - 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 5876.159.102 or higher. This feature is not supported on VMAX systems running HYPERMAX OS 5977 or higher.

To manage VP compression on storage groups:

1. Select the storage 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.
   5. Click **OK**.

Viewing storage groups

Viewing storage groups (HYPERMAX OS 5977 or higher)

This procedure explains how to view storage groups on a storage system running HYPERMAX OS 5977 or higher. There are multiple ways to view the same information. Depending on the method you use, some of the properties and controls may not apply.

1. Select the storage system.

   The **Storage Groups Dashboard** opens by default.
2. Do one of the following:
   - To view *all* storage groups on the storage system, click **Total** in the **SLO Compliance** panel to open the **Storage Groups** list view.

- To view *only* file storage groups (in eNAS environments):
  a. Select **System › Dashboard** to open the **System** dashboard.
  b. In the **Summary** panel, click **File Dashboard** to open the **File Dashboard**.
  c. In the **Summary** panel, click **File Storage Groups** to open the **Storage Groups for File** list view.

The following properties display:

- **Name** — Name of the storage group. Parent storage groups are preceded by an arrow (›) icon. Click the icon to view child storage groups.
- **SLO Compliance** — How well the storage group is complying with its Service Level Objective, if applicable. Possible values are:
  - Critical — Storage group is performing *well* below SLO targets.
  - Marginal — Storage group is performing below SLO target.
  - Stable — Storage group is performing within the SLO target.
  - No SLO — Storage group has no assigned SLO.
  - — Compliance information is being collected.
- **SRP** — Name of SRP that the storage group belongs to, if any.
- **SLO** — Name of the SLO associated with the storage group. If there is no SLO associated with the group, then file displays N/A.
- **Workload Type** — Workload associated with the storage group.
- **Capacity (GB)** — Total capacity of the storage group in GB.
- **Snapshots** — Number of snapshots linked to the storage group.
- **Masking Views** — Number of masking views associated with the storage group.
- **File System #** — Number of file systems associated with the file storage group. This column only displays in eNAS environments.
- **Storage Pools** — Number of storage pools associated with the file storage group. This column only displays in eNAS environments.

The following controls are available:

- **Create SG** — See Creating storage groups (Enginuity 5874 - 5876) on page 120.
- **Modify** — Modifying storage groups on page 129
- **Provision Storage to Host** —Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106/**Provision Storage for File** — Provisioning storage for file on page 805
- **Protect** — Protecting storage groups on page 132

- **Manage Protection** — Managing TimeFinder/Snap sessions on page 396Select a storage group and click **Manage Protection** to display a link dialog box with links to the various replication technologies applied to a storage group—TimeFinder SnapVX, TimeFinder Clone, TimeFinder Mirror, SRDF storage groups, SRDF device groups, and SRDF/Metro groups.

  Note that if a storage group has SRDF Protection with a device group and a storage group, then two management options for SRDF are displayed.
- **View Details** — Viewing storage group details (HYPERMAX OS 5977 or higher) below
- **Move to EMC CloudArray** —Moving storage groups to EMC CloudArray
- **FAST Array Advisor** — Analyzing FAST migration on page 202
- **Change SRP** — Changing Storage Resource Pools for storage groups on page 170
- **Set Host I/O Limits** —Setting host I/O limits on page 146
- **Replication QoS** — See Setting replication QoS on page 232.
- **Delete** —Deleting storage groups on page 145
- **Expand ProtectPoint** — Expanding ProtectPoint storage groups on page 286
- **Start Allocate/Free/Reclaim** — Managing thin pool allocations on page 263/Managing thin pool capacity on page 267/Managing space reclamation on page 265
- **Stop Allocated/Free Reclaim** — Managing thin pool allocations on page 263/Managing thin pool capacity on page 267/Managing space reclamation on page 265
- **Convert to Cascaded** — Converting storage groups to cascaded on page 140
- **Set Volume Status** —Setting volume status on page 229

Viewing storage group details (HYPERMAX OS 5977 or higher)

This procedure explains how to view configuration details for storage groups on storage systems running HYPERMAX OS 5977 or higher. In eNAS operating environments, there are multiple ways to view the same information. Depending on the method you use, some of the properties and controls may not apply.

1. Select the storage system.

   The **Storage Groups Dashboard** opens by default.

2. Do one of the following:

   - To view *all* storage groups on the storage system, click **Total** in the **SLO Compliance** panel to open the **Storage Groups** list view.
   - To view *only* file storage groups (in eNAS environments):

     a. Select **System › Dashboard** to open the **System** dashboard.

     b. In the **Summary** panel, click **File Dashboard** to open the **File Dashboard**.

     c. In the **Summary** panel, click **File Storage Groups** to open the **Storage Groups for File** list view.

3. Select the storage group and click **View Details** to open its details view.

4. If not already displaying, click the **Details** tab.

   The **Storage Group Configuration** view contains Properties, Related Objects, and Performance Views panels. In addition, for file storage groups, the File Systems panel displays. SnapVX, Clone, Mirror and SRDF/Metro also appear in the Related Objects Panel depending on the storage group selection. To manage protection, click the SRDF link in the Associated with section in the Related Objects panel, click **View Storage Groups** to navigate to the SRDF Storage Group list view for the selected storage group, and then click **Manage Protection**. Note that if an SRDF pair is manageable by a storage group and also manageable by a device group, then two management options for SRDF pairs are displayed.

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. If you rename a storage group that has workload on it, you will have to wait some time before the workload is visible in the storage group's **Details** view.
- **SLO Compliance**— How well the storage group is complying with its Service Level Objective, if applicable.
- **SLO** — Service level objective associated with the storage group. If there is no SLO associated with the group, then this field displays N/A.
- **Workload Type** — Type of the workload associated with the storage group.
- **Storage Resource Pool** — Storage resource pool containing the storage group.
- **Total Capacity (GB)** — Total capacity of the storage group in GB.
- **Volumes** — Number of volumes in the storage group.
- **Masking Views** — Number of masking views associated with the storage group.
- **Last Updated** — Timestamp of the most recent changes to the storage group.

- **Host I/O Limit** — Whether the host I/O limit feature is enabled. For more information, see Setting host I/O limits on page 146.
- **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 Enginuity 5876.163.105 or higher. 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.
  - **OnFailure** — 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:

- **Modify** —Modifying storage groups on page 129
- **Provision Storage to Host** —Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106/**Provision Storage for File** — Provisioning storage for file on page 805
- **Delete** — See Deleting storage groups on page 145.
- **Set New Performance Baseline Expectations** — Sets the characteristics currently measured for the previous two weeks as the new performance baseline. See Resetting Workload Plan on page 202.
- **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 **Volumes - 2** will open a view listing the two volumes contained in the the storage group.

Performance Views panel

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

This panel will display with inactive links if the selected Symmetrix system is not registered for data collection.

File Systems panel

The **File Systems** panel displays the following information:

- **File System Name** — File system name.
- **Storage Groups** — Storage groups mapped to the file system.
- **Storage Pools** — Storage pools in the file system.
- **Capacity (GB)** — Total capacity of the file system.

Viewing storage group workload (HYPERMAX OS 5977 or higher)

This procedure explains how to view the statistics used to compute the SLO compliance of a storage group.

Before you begin:

- To perform this operation, a Monitor role is required.
- The storage system must be local and registered for performance.

Viewing storage group workload :

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group and click **View details** to open its details view.
5. If not already displaying, click the **SLO Compliance** tab.

   The **Workload** tab allows you to view workload details, organized into the following panels:

   Service Level

   The service level associated with workload (Storage Group).

   Service Level Compliance

   The calculated compliance of the workload as measured against its SLO.

   Capacity Trend

   The percentage of change in the storage group capacity over the previous two weeks.

   Performance

   Links you to the performance **Analyze**and **Monitor** views for the storage group.

   This panel will display with inactive links if the selected VMAX system is not registered for data collection.

   Weighted Averages

   Displays the following charts:

- **Response Time (ms)** — Displays the calculated weighted average response time in milliseconds for the storage group over the previous 4 hours and 2 weeks, against the SLO value range for the storage group.
- **IOs per Sec** — Displays the calculated weighted average IOs/second for the storage group over the previous 4 hours and 2 weeks.

CPU/Port Load Scores

Displays the following charts:

- **% Front end** — Displays the calculated IO load scores of the storage group's front-end CPUs and ports as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.
- **% Back end** — Displays the calculated IO load scores of the storage group's back-end CPUs and ports as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.

Note You can optionally toggle the format used to display the load scores between **IOs/sec** and **MB/sec**.

Access Density Skew Scores:

Displays the load scores for the storage group in the following charts:

- **% High Load** — Displays the calculated skew density score for the high load disks in the storage group as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.
- **% Mixed Load** — Displays the calculated skew density score for the mixed load disks in the storage group as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.
- **% Low Load** — Displays the calculated skew density score for the low load disks in the storage group as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.

You can toggle the format used to display the load scores between the following views:

- **IOs/sec** — Displays the calculated skew scores for the storage group in IOs per second.
- **MB/sec** — Displays the calculated skew scores for the storage group in megabytes per second.
- **GB** — Displays the calculated skew scores for total data throughput (in GB) for the storage group.

In general, the **High load** and **Low load** numbers should be between 25 and 175 and the **Mixed load** number should be between 50 and 150.

Higher skew storage groups should display the following characteristics for load (**IOs/sec**) and **Capacity**:

- **IOs/sec** — Higher **High load** numbers and lower **Low load** numbers.
- **Capacity** — Higher **Low load** numbers and lower **High load** numbers.

Lower skew storage groups should display higher **Mixed load** numbers and lower **High load** and **Low load** numbers for load (**IOs/sec**) and **Capacity**.

When displaying load (**IOs/sec**):

- Higher **High load** numbers indicate that it is easier for the system to provide lower than expected response time for the storage group.
- Higher **Low load** numbers indicate that it is more difficult for the system to keep the storage group away from SATA storage.

When displaying **Capacity**:

- Lower **High load** numbers indicate that the storage group can have low response times while using less than expected amounts of Flash storage.
- Lower **Low load** numbers indicate that the storage group cannot take advantage of SATA storage.

The following controls are available if the storage system is running HYPERMAX OS 5977 or higher:

- **Save as Reference Workload** — Reference workloads on page 114
- **Reset Workload Plan** — Resetting Workload Plan on page 202

Viewing storage groups (Enginuity 5874 - 5876)

1. Select the storage 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.

   The following properties display:

   - **Name** — Name of the storage group. Parent storage groups are preceded by an arrow (›) icon. Click the icon to view child storage groups.
   - **FAST Policy** — Policy associated with the storage group.
   - **Capacity (GB)** — 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 SG** — Creating storage groups (Enginuity 5874 - 5876) on page 120
   - **Expand** — Expanding storage groups (Enginuity 5874 - 5876) on page 126
   - **Provision Storage to Host** — Using the Provision Storage wizard (Enginuity 5874 - 5876) on page 110
   - **Protect** —Protecting storage groups on page 132

- **Manage Protection** — Managing TimeFinder/Clone sessions on page 379, TimeFinder/Mirror on page 423,Managing TimeFinder/Snap sessions on page 396See "Managing remote replication sessions" on page 441. Select a storage group and click **Manage Protection** to display a link dialog box with links to the various replication technologies applied to a storage group -TF Clone, TF SnapVX (for V3's), TF Mirror, SRDF storage groups, and SRDF device groups.

  Note that if a storage group has SRDF Protection with a device group and a storage group, then two management options for SRDF are displayed.
- **View Details** — Viewing storage group details on the facing page
- **FAST Array Advisor** — Analyzing FAST migration on page 202
- **Delete** — Deleting storage groups on page 145
- **Associate to FAST** — Associating FAST policies and storage groups on page 189
- **Bind SG** — Binding/Unbinding/Rebinding thin volumes on page 273
- **Unbind SG** — Binding/Unbinding/Rebinding thin volumes on page 273
- **Rebind SG** — Binding/Unbinding/Rebinding thin volumes on page 273
- **Reassociate to FAST Policy** — Reassociating FAST polices and storage groups on page 192
- **Disassociate** — Disassociating FAST policies and storage groups on page 191
- **Set Host I/O Limits** —Setting host I/O limits on page 146
- **Set Optimized Read Miss** —Setting optimized read miss on page 227
- **Pin SG** — See Pinning and unpinning volumes on page 197.
- **Unpin SG** — See Pinning and unpinning volumes on page 197.
- **VLUN Migration** — See Migrating regular storage group volumes on page 484.
- **Assign Dynamic Cache Partition** — See Assigning dynamic cache partitions on page 875.
- **VP Compression** —See Managing VP compression on storage groups on page 148.
- **Replication QOS** — See Setting replication QoS on page 232.
- **Start Allocate/Free/Reclaim** — Managing thin pool allocations on page 263/Managing thin pool capacity on page 267/Managing space reclamation on page 265
- **Stop Allocated/Free Reclaim** — Managing thin pool allocations on page 263/Managing thin pool capacity on page 267/Managing space reclamation on page 265
- **Assign Symmetrix Priority** — See Assigning Symmetrix priority on page 223. (Does not apply/display for Symmetrix VMAX 10K (VMAXe) systems)
- **Tag for RecoverPoint** — See Tagging and untagging volumes for RecoverPoint on page 518.

- **Untag for RecoverPoint** — See Tagging and untagging volumes for RecoverPoint on page 518.

Viewing storage group details

1. Select the storage 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. Clone and Mirror also appear in the Related Objects Panel depending on the storage group selection. To manage protection, click the SRDF link in the Associated with section in the Related Objects panel, click **View Storage Groups** to navigate to the SRDF Storage Group list view for the selected storage group, and then click **Manage Protection**. Note that if an SRDF pair is manageable by a storage group and also manageable by a device group, then two management options for SRDF pairs are displayed.

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 (GB)** — Total capacity of the storage group in GB.
- **Volumes** — Number of volumes in the storage group.
- **Masking Views** — Number of masking views associated with the storage group.
- **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 146.
- **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 Enginuity 5876.163.105 or higher. 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.
  - **OnFailure** — 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:

- **Expand** — See Expanding storage groups (Enginuity 5874 - 5876) on page 126.
- **Delete** — See Deleting storage groups on page 145.
- **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 will display with inactive links if the selected storage 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.

Viewing backup storage groups (HYPERMAX OS 5977 or higher)

This procedure explains how to view storage groups protected by ProtectPoint.

1. Select the storage system.
2. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard**.
3. In the **Storage Group Protection** panel, click **Backup** to open the **Backup Replication** view.

   The following properties display:

   Note Partially protected storage groups will not appear in this view. For example, after adding a volume to a storage group protected by ProtectPoint, the storage group is considered partially protected, in which case it will no longer appear in this view. The storage group will remain in this state until you expand it using ProtectPoint, after which it will be considered fully protected, and reappear in this view.

   - **State** — Replication state severity of the storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

     Normal — All volumes are in a normal state.

     Warning — One or more volumes in the group are in a warning state, with none higher.

     Error — One or more volumes in the group are in an Error state.
   - **Storage Group** — Name of the storage group.
   - **Snapshot** — Name of the snaphshot.
   - **Capacity** — Total capacity of the storage group in GB.
   - **Backup Storage Group** — Name of the backup storage group.
   - **Last Creation Time** — Date, time, and generation number for the snapshot.

The following control is available:

- **Manage Protection** — Viewing snapshots on page 417

Viewing point-in-time storage groups

This procedure explains how to view storage groups protected by TimeFinder/SnapVX, TimeFinder/Clone, or TimeFinder/Mirror.

1. Select the storage system.
2. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard** (applies only for storage systems running HYPERMAX OS 5977 or higher).
3. In the **Storage Group Protection** panel, click **Point in Time** to open the **Point in Time** view.

The **Point in Time** view consists of the following tabs, depending on the storage environment:

SnapVX:

Applies only for storage systems running HYPERMAX OS 5977 or higher.

The following properties display:

- **State** — Replication state severity of the storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

  Normal — All volumes are in a normal state.

  Warning — One or more volumes in the group are in a warning state, with none higher.

  Error — One or more volumes in the group are in an Error state.
- **Storage Group** — Name of the storage group.
- **Capacity (GB)** — Total capacity of the storage group in GB.
- **Snapshot** — Name of the snaphshot.
- **Last Creation Time** — Date, time, and generation number for the snapshot.

The following controls are available:

- **Create Snapshot** — Creating snapshots on page 409
- **Link** — Linking to snapshots on page 410
- **Restore** — Restoring snapshots on page 413
- **Unlink** — Unlinking from snapshots on page 413
- **Manage Protection** — Viewing snapshots on page 417

Clone:

- **State** — Replication state severity of the storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

  Normal — All volumes are in a normal state.

  Warning — One or more volumes in the group are in a warning state, with none higher.

Error — One or more volumes in the group are in an Error state.

- **Storage Group** — Name of the storage group.
- **Replication Management Group** — Replication management group containing the storage group.
- **Link States** — Lists the pair states found in the management group.
- **Group Type** — Type of management group.

The following controls are available:

- **Activate** —See Activating clone copy sessions on page 382.
- **Recreate** —See Recreating clone copy sessions on page 383.
- **Split** —See Splitting clone volume pairs on page 388.
- **Restore** —See Restoring data from target volumes on page 387.
- **Manage Protection** — Viewing snapshots on page 417

Mirror:

- **State** — Replication state severity of the storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

  Normal — All volumes are in a normal state.

  Warning — One or more volumes in the group are in a warning state, with none higher.

  Error — One or more volumes in the group are in an Error state.
- **Storage Group** — Name of the storage group.
- **Replication Management Group** — Replication management group containing the storage group.
- **Link States** — Lists the pair states found in the management group.
- **Group Type** — Type of management group.

The following controls are available:

- **Establish** —See Creating Snapshots on page 424.
- **Restore** —See Restoring BCV pairs on page 425.
- **Split** —See Splitting BCV pairs on page 427.
- **Cancel** —See Cancelling BCV pairs on page 427.
- **Manage Protection** — Viewing snapshots on page 417

Viewing remote replication storage groups

This procedure explains how to view storage groups protected by SRDF. Note that for the Protection Dashboard showing SRDF protected SGs, only R1 SGs will appear in the Application Protection section and Storage Group Protection counts. An R2 SG will appear in the SG Protection count and Application Protection list, if it is an R1 in an R21 setup.

1. Select the storage system.
2. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard** (applies only for storage systems running HYPERMAX OS 5977 or higher).
3. In the **Storage Group Protection** panel, click **Remote Replication** to open the **Remote Replication** view.

   The following properties display:

   - **State** — Replication state severity of the storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

     Normal — All pairs are in a normal state.

     Warning — One or more oairs in the group are in a warning state, with none higher.

     Error — One or more pairs in the group are in an Error state.
   - **Storage Group** — Name of the storage group.
   - **Replication Management Group** — Replication management group associated with the storage group.
   - **Link States** — State of the SRDF links.
   - **Multiple Replicas** — SRDF configuration.

   The following controls are available:

   - **Establish** — See Establishing SRDF pairs on page 461.
   - **Split** — See Splitting SRDF pairs on page 477.
   - **Suspend** — Suspending SRDF pairs on page 478
   - **Restore** — See Restoring SRDF pairs on page 474.
   - **Resume** — Resuming SRDF links (page 470)
   - **Failover** — Failing over (page 462)
   - **Failback** — Failing back (page 463)
   - **Manage Protection**— Select a Storage Group and click **Manage Protection**— to navigate to the SRDF Storage Group List view displaying the SRDF storage groups.

Viewing high availability storage groups (HYPERMAX OS 5977 or higher)

This procedure explains how to view high availability SRDF/metro storage groups.

1. Select the storage system.

2. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard** (applies only for storage systems running HYPERMAX OS 5977 or higher).
3. In the **Storage Group Protection** panel, click **High Availability** to open the **Remote Replication** view.

   The following properties display:

   - **State** — State severity of the SRDF/Metro storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

     Normal — All pairs are in a normal state.

     Warning — One or more oairs in the group are in a warning state, with none higher.

     Error — One or more pairs in the group are in an Error state.
   - **Storage Group** — Name of the storage group.
   - **Replication Management Group** — Replication management group associated with the storage group.
   - **Link States** — State of the SRDF links.
   - **Bias Location** — Bias location.

   The following controls are available:

   - **Establish** — See Establishing SRDF pairs on page 461.
   - **Suspend** — Suspending SRDF pairs on page 478
   - **Restore** — See Restoring SRDF pairs on page 474.
   - **Set Bias Location**— Setting bias location (page 475)
   - **Manage Protection**— Select a Storage Group and click **Manage Protection**— to navigate to the SRDF/Metro Storage Group List view displaying the SRDF/Metro storage groups.

Viewing total storage groups

This procedure explains how to view all storage groups on a storage system, regardless of whether they are proctected:

1. Select the storage system.
2. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard** (applies only for storage systems running HYPERMAX OS 5977 or higher).
3. In the **Storage Group Protection** panel, click **Total**to open the **All Storage Groups** list view.

The following properties display:

- **Storage Group** — Name of the storage group. Parent storage groups are preceded by an arrow (›) icon. Click the icon to view child storage groups.
- **Capacity (GB)**— Total capacity of the storage group in GB.
- **SLO** — Service Level Objective associated with the storage group.
- **Masking Views** — Number of masking views associated with the storage group.

The following controls are available:

- **Protect** — See Protecting storage groups on page 132.
- **Manage Storage Groups** — Managing storage groups (HYPERMAX OS 5977 or higher) on page 117
- **Set SRDF GCM** — Setting the SRDF GCM flag on page 475
- **Reset SRDF/Metro Identity**— Resetting original device identity on page 473
- **Set Volume Status**— Setting volume status on page 476

Viewing unprotected storage groups

This procedure explains how to view unprotected storage groups.

1. Select the storage system.
2. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard**.
3. In the **Storage Group Protection** panel, click **Unprotected** to open the **Unprotected Storage Groups** list view.

   The following properties display:

   - **Storage Group** — Name of the storage group. Parent storage groups are preceded by an arrow (›) icon. Click the icon to view child storage groups.
   - **Capacity (GB)** — Total capacity of the storage group in GB.
   - **Masking Views** — Number of masking views associated with the storage group.

   The following controls are available:

   - **Protect** — See Protecting storage groups on page 132.
   - **Manage Storage Groups** — Managing storage groups (HYPERMAX OS 5977 or higher) on page 117

Fully Automated Storage Tiering

Understanding FAST (HYPERMAX OS 5977)

Note This section describes FAST operations for storage systems running HYPERMAX OS 5977. Understanding FAST (Enginuity 5874 - 5876) on page 176 describes FAST operations on storage systems running Enginuity 5874 - 5876.

Fully Automated Storage Tiering (FAST) automates management of storage system disk resources on behalf of thin volumes. FAST automatically configures disk groups to form a Storage Resource Pool by creating thin pools according to each individual disk technology, capacity and RAID type.

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

FAST further provides the ability to deliver variable performance levels through Service Level Objectives. Thin volumes can be added to storage groups and the storage group can be associated with a specific service level objective to set performance expectations.

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

Note On storage systems running HYPERMAX OS 5977 or higher, FAST is always on. The *VMAX3 Family Product Guide* provides additional details about FAST operations in this storage environment.

Understanding Service Level Objectives

A service level objective is the response time target for the storage group. The service level objective allows you set the VMAX array with the desired response time target for the storage group. It automatically monitors and adapts to the workload in order to maintain (or meet) the response time target. The service level objective includes an optional workload type so you can further tune expectations for the workload storage group to provide just enough flash to meet your performance objective.

Renaming Service Level Objectives

This procedure explains how to rename Service Level Objectives. Once renamed, all active management and reporting activities will be performed on the newly named Service Level Objective. The original, pre-configured Service Level Objective name will be maintained in the Service Level View for future reference. All other references to the original Service Level Objective will display the new name.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- This feature requires HYPERMAX OS 5977 or higher.

- Service Level Objective must be unique from other Service Level Objectives on the storage system and cannot exceed 32 characters. Only alphanumeric characters, underscores (_), and hyphens (-) are allowed. However, Service Level Objective names cannot start or end with an underscore or hyphen.

To rename Service Level Objectives:

1. Select the storage system.
2. Select **Storage › Service Levels** to open the **Service Level** view.
3. Hover the cursor over the Service Level Objective and click Edit ( ).
4. Type the new name over the existing name.
5. Click **OK**.

Reverting to original Service Level Objective names

This procedure explains how to revert a renamed Service Level Objective to its original, pre-configured name.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- This feature requires HYPERMAX OS 5977 or higher.
- Service Level Objective must be unique from other Service Level Objectives on the storage system and cannot exceed 32 characters. Only alphanumeric characters, underscores (_), and hyphens (-) are allowed. However, Service Level Objective names cannot start or end with an underscore or hyphen.

To rename Service Level Objectives:

1. Select the storage system.
2. Select **Storage › Service Levels** to open the **Service Level** view.
3. Hover the cursor over the Service Level Objective and click Edit ( ).
4. Type the original, pre-configured name.
5. Click **OK**.

Viewing Service Level Objectives

This procedure explains how to view information on the available Service Level Objectives.

A service level objective is the response time target for the storage group. The service level objective allows you set the VMAX array with the desired response time target for the storage group. It automatically monitors and adapts to the workload in order to maintain (or meet) the response time target. The service level objective includes an optional workload type so you can further tune expectations for the workload storage group to provide just enough flash to meet your performance objective.

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view Service Level Objectives:

1. Select the storage system.
2. Select **Storage › Service Levels** to open the **Service Level Objectives** view.

   Available Service Level Objectives appear in white and unavailable appear grayed-out.
3. *Optional*: To rename a Storage Level Objective, hover the cursor over the Service Level Objective and click Edit ( ). Type the new name over the existing name.
4. *Optional*: To provision storage using a Storage Level Objective, hover the cursor over the Storage Level Objective and click **Provision Storage**. This will open the Provision Storage to Host wizard, with the Service Level Objective and Storage Resource Pool already selected. For more information on using the wizard, refer to its help page.
5. To view details on the **Expected Average Response Times**, click **Show Details**. Response time details for each Service Level Objective display. Where OLTP workload is focused on optimizing performance for small block I/O and DSS workload is focused on optimizing performance for large block I/O. The Workload Type can also specify whether to account for any overhead associated with replication (OLTP_Rep and DSS_Rep, where denotes replication).

Understanding Storage Resource Pool details

Storage Resource Pool is a collection of data pools that provide FAST a domain for capacity and performance management. By default, a single default Storage Resource Pool is factory pre-configured. Additional Storage Resource Pools can be created with a service engagement. FAST performs all its data movements within the boundaries of the Storage Resource Pool.

Modifying Storage Resource Pool details

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- You must have Administrator or StorageAdmin permission.

To modify Storage Resource Pools:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** panel, under **Reports**, click **View SRP Details** to open the **SRP Details** view.
4. Modify any number of the following:
   - **Name** — Name of the storage resource pool. To change this value, type a new description and click **Apply**.
     The name of the storage resource pool must be unique and it cannot exceed 32 characters. It can include only alphanumeric, underscore, and hyphen characters, but cannot begin with an underscore or hyphen character.

- **Description** — Optional description of the pool. To change this value, type a new description and click **Apply**.
  The description cannot exceed 127 characters. It can contain only alphanumeric, hyphen, underscore, space, period, and comma characters.
- **Reserved Capacity % (0 - 80)** The percentage of the capacity of the Storage Resource Pool to be reserved for volume write I/O activities. Valid values for the percentage are from 1 to 80. NONE disables it.

  For example, if you set the reserved capacity on a Storage Resource Pool to 30%, then the first 70% of the pool capacity is available for general purpose operations (host I/O allocations, local replication tracks and SRDF/A DSE allocations) and the final 30% of the pool capacity is reserved strictly for volume write I/O activities.

  Note that existing TimeFinder snapshot sessions created on volumes in the Storage Resource Pool are invalid if the free capacity of the Storage Resource Pool, as a percentage of the usable capacity, goes below the reserved capacity.
- **Usable by RDFA DSE** — Specifies whether the Storage Resource Pool can be used for SRDF/A DSE operations. This field does not display for external SRPs.
  The maximum amount of storage from a Storage Resource Pool that can be used for DSE is controlled by the system wide dse_max_cap setting, as described in the *Solutions Enabler SRDF Family CLI User Guide*.

5. Click **Apply**.

Viewing Storage Resource Pools

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- This procedure explains how to open the **Storage Resource Pool** view from the main menu. You can also open it from **Storage Groups Dashboard**.

To view Storage Resource Pools:

1. Select the storage system.
2. Select **Storage › Storage Resource Pools** to open the **Storage Resource Pools** view.
3. Select the tab for the storage resource pool you want to view.

   The Storage Resource Pool view allows you to view and manage storage resource pools. It contains Properties, Related Objects, and Disk Groups panel: on page 170 panels.

Properties panel:

The following properties display:

- **Name** — Name of the storage resource pool. To change this value, type a new description and click **Apply**.
  The name of the storage resource pool must be unique and it cannot exceed 32 characters. It can include only alphanumeric, underscore, and hyphen characters, but cannot begin with an underscore or hyphen character.
- **Description** — Optional description of the pool. To change this value, type a new description and click **Apply**.
  The description cannot exceed 127 characters. It can contain only alphanumeric, hyphen, underscore, space, period, and comma characters.
- **Default Emulation** — The default emulation for the pool (FBA or CKD).
- **Usable Capacity (GB)** Usable capacity of all the disk groups in the Storage Resource Pool, excluding any external disk groups used for FTS encapsulation.
- **Allocated Capacity (GB)** — Sum of the volume allocations, snapshot allocations, and SRDF/A DSE allocations on the Storage Resource Pool.
- **Free Capacity (GB)** — Difference between the usable and allocated capacities.
- **Subscription (GB)** — Percentage of the configured sizes of all the thin volumes subscribed against the Storage Resource Pool.
- **Reserved Capacity % (0 - 80)** — Percentage of the **Usable Capacity** that will be reserved for non-snapshot activities. Existing TimeFinder snapshot sessions created on volumes in the Storage Resource Pool can go invalid if the Free Capacity of the Storage Resource Pool, as a percentage of the Usable Capacity, goes below the Reserved Capacity. To change this value, type a new value over the old, and click **Apply**.
- **Usable by RDFA DSE** — Specifies whether the Storage Resource Pool can be used for SRDF/A DSE operations. This field does not display for external SRPs.

The following controls are available:

- **Apply** — Applies changes made in the **Properties** list. For example, changing the Reserved Capacity %.
- **Cancel** — Cancels changes made in the **Properties** list.
- **Add eDisk** — Adding external disks on page 290

Related Objects panel:

The **Related Objects** panel links you to views displaying objects contained in the pool. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Service Levels - 2** will open a view listing the two service levels in the storage resource pool.

Disk Groups panel:

The **Disk Groups** panel lists details on the disk groups in the SRP, including:

- **Name** — Name of the disk group.
  For external disk groups, clicking on the disk group name opens a list of external disks. For more information, refer to Filtered results table.
- **Vendor** — Vendor and product ID.
- **Speed (RPM)** — Speed of the disks in the disk group.
- **Usable Capacity (GB)** —Usable capacity in the disk group.
- **Disk Location** — Specifies if the disk group is internal or external.

Changing Storage Resource Pools for storage groups

This procedure explains how to change the Storage Resource Pool of a parent storage group, with child SLOs using different Storage Resource Pools.

This procedure explains how to perform this operation from the **Storage Groups Dashboard**. In eNAS environments, you can also perform this operation from the **File Storage Groups** page (**System › Dashboard › File Dashboard › File Storage Groups**).

Before you begin:

- The storage system must be running HYPERMAX OS 5977 or higher.
- You must have Administrator or StorageAdmin permission.

To change the Storage Resource Pool for storage groups:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group, click more >> , and select **Change SRP** to open the **Change SRP** dialog box.
5. Select the new **SRP**.
6. *Optional*: Change the **Service Level** for the SG. Service level objectives specify the characteristics of the provisioned storage, including maximum response time, workload type, and priority. This field defaults to None if you set the **Storage Resource Pool** to None. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

7. *Optional*: Refine the SLO by selecting the **Workload Type**to assign it.
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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing Storage Resource Pool alerts

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view Storage Resource Pool alerts:

1. Select the storage system.
2. Select **Storage › Storage Group Dashboard** to open the **Storage Group Dashboard**.
3. In the **Storage Resource Pools** panel, click the number of SRP alerts 0 to open the **SRP Alerts** list 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.
- **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** — Viewing Storage Resource Pool alert details below
- **Acknowledge** — Acknowledging Storage Resource Pool alerts below
- **Delete** — Deleting Storage Resource Pool alerts on the facing page

Viewing Storage Resource Pool alert details

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view Storage Resource Pool alerts:

1. Select the storage system.
2. Select **Storage › Storage Group Dashboard** to open the **Storage Group Dashboard**.
3. In the **Storage Resource Pools** panel, click the number of SRP alerts 0 to open the **SRP Alerts** list view.
4. Select the alert and click **View Details**.

Acknowledging Storage Resource Pool alerts

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view Storage Resource Pool alerts:

1. Select the storage system.
2. Select **Storage › Storage Group Dashboard** to open the **Storage Group Dashboard**.
3. In the **Storage Resource Pools** panel, click the number of SRP alerts 0 to open the **SRP Alerts** list view.
4. Select one or more alerts and click **Acknowledge**.

Deleting Storage Resource Pool alerts

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view Storage Resource Pool alerts:

1. Select the storage system.
2. Select **Storage › Storage Group Dashboard** to open the **Storage Group Dashboard**.
3. In the **Storage Resource Pools** panel, click the number of SRP alerts 0 to open the **SRP Alerts** list view.
4. Select one or more alerts and click **Delete**.

SLO Compliance

Each Service Level Objective and workload type has a response band associated with it. When a storage group (workload) is said to be compliant, it means that it is operating within the associated response time band.

When assessing the compliance of a storage group, Workload Planner calculates its weighted response time for the past 4 hours and for the past 2 weeks, and then compares the two values to the maximum response time associated with its given SLO. If both calculated values fall within (under) the SLO defined response time band, the compliance state is STABLE. If one of them is in compliance and the other is out of compliance, then the compliance state is MARGINAL. If both are out of compliance, then the compliance state is CRITICAL.

Creating SLO Compliance Reports

This procedure explains how to create SLO Compliance reports. SLO Compliance Reports allow you to view storage group performance against Service Level Objectives over a period of time.

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- The user must be a StorageAdmin permissions or higher.

To create SLO Compliance Reports:

1. Select the storage system.

2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Group Management** panel, under **Manage**, click **View SLO Compliance Report** to open the report.
4. Click **Schedule** to open the **Create Report** dialog box.
5. On the **General** tab, do any number of the following:
   a. Type a **Name** for the report.
   b. Type a **Description** for the report.
   c. Select the time zone in which the report will be generated (**Generated Time Zone**).
6. On the **Schedule** tab, do any number of the following:
   a. Select the **First Runtime**.
   b. Select the **Day(s) to Run**.
   c. Select the number of days that the report should be retained (**Retention (Days)**).
7. On the **Email** tab, select **Send report to** and type an email address.
8. Click **OK**.

Viewing SLO Compliance Reports

This procedure explains how to view storage group performance against Service Level Objectives over a period of time.

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- The user must be a StorageAdmin permissions or higher.

To view SLO compliance reports:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Group Management** panel, under **Manage**, click **View SLO Compliance Report** to open the report.
4. *Optional*: Customize the report by doing the following:
   a. Select to view **All Service Levels** or a specific service level.
   b. Select the time period. For the time period you select, WLP will assess the storage group's compliance in 30 minute intervals, and then display its overall compliance state based on the method described in SLO Compliance on the previous page. For example, if you select **Last 24 hours**, WLP will assess the storage group's compliance state 48 times, and then display its calculated compliance state in this report.
   c. Select whether to view the compliance information as a chart or numerical display.

   The following properties display:

   - **Storage Group** — Name of the storage group.
   - **Service Level** — Service level associated with the storage group.

- **% Stable** — Percentage of time the storage group performed within the SLO target.
- **% Marginal** — Percentage of time the storage group performed below the SLO target.
- **% Critical** — Percentage of time the storage group performed *well* below the SLO target.

The following controls are available:

- **Save As** —Save the report as a JPG, PNG, or PDF file
- **Schedule** —Exporting the Storage Group Demand Reports below
- **Monitor** — Open the Performance Monitor page for the storage group

Storage Group Demand Reports

Viewing Storage Group Demand Reports

This procedure explains how to view storage groups on an SRP and their associate workloads.

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view storage group demand reports:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** panel, under **Reports**, click **Storage Group Demand Report** to open the report.

   The following properties display:

   - **Storage Group** — Name of the storage group.
   - **Subscription (GB)** — Amount of SRP capacity to which the storage group subscribed.
   - **Allocated (%)** — The percentage of allocated pool capacity.
   - **Snap Allocated (%)** — The percentage allocated to snap.

   The following control is available:

   - **Export Report** — Exporting the Storage Group Demand Reports below

Exporting the Storage Group Demand Reports

This procedure explains how to export a portable document format (PDF) version of the Storage Group Demand Report:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** panel, under **Reports**, click **Storage Group Demand Report** to open the report.
4. Click **Export Report**.

5. Click **OK** in the confirmation message.
6. Select a location for the download and click **Save**.

Understanding FAST (Enginuity 5874 - 5876)

Note This section describes FAST operations for storage systems running Enginuity 5874 - 5876. Understanding FAST (HYPERMAX OS 5977) on page 165 describes FAST operations on storage systems running HYPERMAX OS 5977 or higher.

Fully Automated Storage Tiering (FAST) is software that runs background algorithms to continuously analyze the utilization (busy rate) of the storage 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.

Enginutiy 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.

Enginutiy 5876.159.102 or higher 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.

Note On storage systems running Enginuity 5874 - 5876, FAST is a license-able feature. The *Solutions Enabler Array Controls CLI Guide* provides additional details about FAST in this storage environment.

FAST versions

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

Table 2: Comparison of FAST and Fast Virtual Pools (VP)

FAST DP	FAST Virtual Pools
Requires Enginuity 5874 - 5876	Requires Enginuity 5875 - 5876
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

FAST DP	FAST Virtual Pools
	512 volumes
Disk group provisioing (DP) tiers: contain disk groups	Virtual pool (VP) tiers: contain thin 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 5876.159.102 or higher

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:



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 195).

Monitoring FAST

Once you have configured FAST on a storage system, you can use the FAST dashboard to monitor its status, as described in Monitoring FAST (Enginuity 5874 - 5876) (page 179).

Monitoring FAST (Enginuity 5874 - 5876)

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.

Before you begin:

- The FAST dashboard is only available for storage systems running Enginuity 5874 - 5876.
- To refresh the information displayed in the **FAST** dashboard, click refresh Symmetrix ( ) in the status bar, or exit/return to the **FAST** dashboard.

To access the FAST dashboard:

1. Select the storage system.
2. Select **Storage › FAST** .

FAST dashboard

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

Fast Type

If the storage system is licensed for both FAST DP 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:

- **Settings** — Opens the **FAST Settings** dialog box, from which you modify FAST settings. For more information, refer to FASTcontroller on page 195.
  - **State** — The current state of FAST on the storage 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. This state does not apply in
    - **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 Approved** — 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.

◆ **Time Windows**

- **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** — Indicates that FAST is operating within a move time window.
  - **Open** — 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 Groups** — Number of storage groups associated with the policy.

Tiers Demand Report

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

The chart format provides the following information for each tier, displayed in **Alphabetical** or **Used** order.:

- **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 to FAST** — The amount of storage available to FAST VP on the tier, in GB. This is the calculated by taking the total enabled capacity for the all the pools in the tier (sum of all enabled TDATs) minus the pool reserved capacity (PRC).

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 (GB)** — 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 Used 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. indicates compliance. indicates non-compliance.

Symmetrix tiers

Creating Symmetrix tiers

This procedure explains how to create Symmetrix tiers.

Before you begin:

- This feature requires Enginuity 5874 - 5876.
- The maximum number of tiers that can be defined on a storage system is 256.
- When a disk group or thin pool is specified, its technology type must match the tier technology.
- Disk groups can only be specified when the tier include type is static.
- A standard tier cannot be created if it 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 storage system.
2. Select **Storage › Tiers** to open the **Tiers** list view.
3. Click **Create Tier**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 **Include all future disk groups on matching 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage system.
2. Select **Storage** › **Tiers** to open the **Tiers** list view.
3. Select the tier and click **Modify** 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 Diskgroup Tiers (Propagate)**.
6. Click **OK**.

Renaming Symmetrix tiers

This procedure explains how to rename Symmetrix tiers.

Before you begin:

- This feature requires Enginuity 5874 - 5876.
- 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 storage 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage system.
2. Select **Storage › Tiers** to open the **Tiers** list view.
3. Select the tier and click **Delete**.
4. Click **Delete** in the confirmation message.

Viewing Symmetrix tiers

Before you begin:

This feature requires Enginuity 5874 - 5876.

To view Symmetrix tiers:

1. Select the storage 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:
  - **Disk Group** — 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.
  - **Virtual Pool** — 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.
- **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 (GB)** — 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 Tier** —See Symmetrix tiers on page 181.
- **Modify** —See Modifying Symmetrix tiers on page 183.
- **View Details** —See Viewing Symmetrix tier details below.
- **Delete** —See Deleting Symmetrix tiers on the previous page.

Viewing Symmetrix tier details

Before you begin:

This feature requires Enginuity 5874 - 5876.

To view Symmetrix tiers details:

1. Select the storage 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 page 187, Performance Views, and Graphs on page 187 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 (GB)** — Amount of free/unused storage on the tier, in GB.

- **Free Capacity (GB)** — 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 (GB)** — Sum of hypers of all volumes in FAST storage group with matching RAID protection that reside on this tier.

- **FAST Free (GB)** — 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 (GB)** —The calculated upper limit for the storage group on the tier.

- **Excess (GB)** — 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.

The following controls are available:

- **Modify** —See Modifying Symmetrix tiers on page 183.
- **Delete** —See Deleting Symmetrix tiers on page 184.
- **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 Views panel

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

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:

- This feature requires Enginuity 5874 - 5876.
- The maximum number of policies allowed per storage 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 contains 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 Array Controls CLI 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 storage 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage 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 **Delete** in the confirmation message.

Associating 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.

Before you begin:

- This feature requires Enginuity 5874 - 5876.
- Storage groups and FAST policies can only be associated under the following conditions:

- The VMAX system is running Enginuity 5874 through 5876.
- 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 storage 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** panel, click **Manage Policies** to open the **Manage Policies** view.
5. Select the policy and click **Associate Storage Groups** to open the **Associate Storage Groups** dialog box.
6. Select one or more storage groups, and then 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 storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group, click more >> and select **Associate to FAST** to open the **Associate to FAST Policy** dialog box.
4. Select a policy and click **OK**.

Disassociating FAST policies and storage groups

1. Select the storage 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 - nn** 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group, click more >>, 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

Before you begin:

This feature requires Enginuity 5874 - 5876.

To view FAST policies:

1. Select the storage 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 Groups** — Storage group associated with the policy.

   The following controls are available:

   - **Create** — See FAST policies on page 187.
   - **View Details** — See Viewing FAST policy details below.
   - **Delete** — See Deleting FAST policies on page 189.
   - **Associate Storage Groups**— Associating FAST policies and storage groups on page 189

Viewing FAST policy details

Before you begin:

This feature requires Enginuity 5874 - 5876.

To view FAST policy details:

1. Select the storage 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** — Associating FAST policies and storage groups on page 189
- **Delete** — See Deleting storage groups on page 145.
- **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 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

Before you begin:

This feature requires Enginuity 5874 - 5876.

To set FAST control parameters:

1. Select the storage system.
2. Select **Storage › FAST** to open the **FAST** dashboard.
3. If the storage system is licensed for both FAST DP and FAST VP, select the **FAST Type**.
4. In the **FAST Status Report** view block, click **Settings** to open the corresponding **FAST [VP|DP] 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 3: Comparison of FAST DP and Fast VP

FAST DP	FAST VP
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.	
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.	
Maximum Moves Per Day (2-200) — 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.	**Relocation Rate (1-10)** — 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.

FAST DP	FAST VP
Maximum Simultaneous Moves (2-32) — 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 option is shared with Symmetrix Optimizer.	**Pool Reserved Capacity (1-80)** — 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 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.
Migration Restriction — Specifies whether the FAST controller can perform swaps and moves, or only swaps.	**Allocate by FAST Policy** — When enabled, the system chooses a pool for the policy when making an allocation for a thin volume.
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.	**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: ◆ **Time to Compress (40-400)** — Specifies how frequently compression will occur. Valid values are 1 - 365 days. Default is 40 days. ◆ **Compression Rate (1-10)** — Specifies the compression rate. Valid values are 1 (most aggressive to 10 (least aggressive). Default is 5. You can enable this feature for a thin pool while creating it (Thin pools (page 257) or post creation in the pool's **Details** view (Viewing thin pool details (page 269)). In addition, you can also manually control this feature at the storage group level (Managing VP compression on storage groups (page 148)).

Pinning and unpinning volumes

Pinning volumes prevents any automated process such as FAST or Optimizer from moving them. However, you can still can 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.

Before you begin:

This feature requires Enginuity 5874 - 5876.

To pin and unpin volumes:

1. Select the storage system.
2. Select **Storage** › **Volumes** to open the **Volumes** dashboard.
3. Select the volume type in tree menu and click **View** to open its list view.
4. Select one or more volumes, click more >> , and select one of the following:
   - **Pin** — To pin the volumes.
   - **Unpin** — To unpin the volumes.
5. 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 Unisphere for VMAX supports time windows in the enhanced format introduced in Enginuity 5875.139.93. To use time windows created in earlier versions of Enginuity,you must convert them according to the procedure in the *Solutions Enabler Array Controls CLI 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:

- This feature requires Enginuity 5874 - 5876.
- 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 storage 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 the type of time window you want to create/modify.

   Depending on your selection, either the **Performance Time Window** or the **Move 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 Windows (Inclusive)** option to open **Manage Open Performance 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 (Fri:18:00 - Mon:00:00)** — Creates a single open time window for the weekend (17:00 Friday to 8:00 Monday).
      - **9:00-17:00 , Monday-Friday** — Creates five time windows, one for each day of the work week.
      - **17:00-8:00, Monday-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 Windows (Exclusive)** option to open **Manage Closed Performance Time Windows** 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 one week.
- **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, with the default being eight hours.

9. Click **OK**.

Understanding Workload Planner

Workload Planner is a FAST component used to display performance metrics for applications and to model the impact of migrating the workload from one storage system to another (Analyzing FAST migration on page 202).

Workload Planner is supported on storage systems running Enginuity 5874 - 5876, or HYPERMAX OS 5977.

For storage groups to be eligible for Workload Planning, they must meet the following criteria:

- On a locally attached storage system registered for performance. See Registering storage systems on page 630 for instructions on registered storage systems.
- Belong to only one masking view.
- Under FAST control:
  - For storage systems running HYPERMAX OS 5977, they must be associated with a Service Level Objective.
  - For storage systems running Enginuity 5874 - 5876, they must be associated with a FAST policy.
- Contain only FBA volumes.

In addition, the Unisphere for VMAX server must be on an open systems host.

Viewing storage group workload (HYPERMAX OS 5977 or higher)

This procedure explains how to view the statistics used to compute the SLO compliance of a storage group.

Before you begin:

- To perform this operation, a Monitor role is required.
- The storage system must be local and registered for performance.

Viewing storage group workload :

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group and click **View details** to open its details view.
5. If not already displaying, click the **SLO Compliance** tab.

   The **Workload** tab allows you to view workload details, organized into the following panels:

   Service Level

   The service level associated with workload (Storage Group).

   Service Level Compliance

   The calculated compliance of the workload as measured against its SLO.

   Capacity Trend

   The percentage of change in the storage group capacity over the previous two weeks.

   Performance

   Links you to the performance **Analyze**and **Monitor** views for the storage group.

   This panel will display with inactive links if the selected VMAX system is not registered for data collection.

   Weighted Averages

   Displays the following charts:

   - **Response Time (ms)** — Displays the calculated weighted average response time in milliseconds for the storage group over the previous 4 hours and 2 weeks, against the SLO value range for the storage group.
   - **IOs per Sec** — Displays the calculated weighted average IOs/second for the storage group over the previous 4 hours and 2 weeks.

   CPU/Port Load Scores

   Displays the following charts:

   - **% Front end** — Displays the calculated IO load scores of the storage group's front-end CPUs and ports as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.
   - **% Back end** — Displays the calculated IO load scores of the storage group's back-end CPUs and ports as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.

   Note You can optionally toggle the format used to display the load scores between **IOs/sec** and **MB/sec**.

Access Density Skew Scores:

Displays the load scores for the storage group in the following charts:

- **% High Load** — Displays the calculated skew density score for the high load disks in the storage group as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.
- **% Mixed Load** — Displays the calculated skew density score for the mixed load disks in the storage group as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.
- **% Low Load** — Displays the calculated skew density score for the low load disks in the storage group as a percentage of the storage group's expected values. The scores are calculated for the previous 4 hours and 2 weeks.

You can toggle the format used to display the load scores between the following views:

- **IOs/sec** — Displays the calculated skew scores for the storage group in IOs per second.
- **MB/sec** — Displays the calculated skew scores for the storage group in megabytes per second.
- **GB** — Displays the calculated skew scores for total data throughput (in GB) for the storage group.

In general, the **High load** and **Low load** numbers should be between 25 and 175 and the **Mixed load** number should be between 50 and 150.

Higher skew storage groups should display the following characteristics for load (**IOs/sec**) and **Capacity**:

- **IOs/sec** — Higher **High load** numbers and lower **Low load** numbers.
- **Capacity** — Higher **Low load** numbers and lower **High load** numbers.

Lower skew storage groups should display higher **Mixed load** numbers and lower **High load** and **Low load** numbers for load (**IOs/sec**) and **Capacity**.

When displaying load (**IOs/sec**):

- Higher **High load** numbers indicate that it is easier for the system to provide lower than expected response time for the storage group.
- Higher **Low load** numbers indicate that it is more difficult for the system to keep the storage group away from SATA storage.

When displaying **Capacity**:

- Lower **High load** numbers indicate that the storage group can have low response times while using less than expected amounts of Flash storage.
- Lower **Low load** numbers indicate that the storage group cannot take advantage of SATA storage.

The following controls are available if the storage system is running HYPERMAX OS 5977 or higher:

- **Save as Reference Workload** — Reference workloads on page 114

- **Reset Workload Plan** — Resetting Workload Plan below

Resetting Workload Plan

This procedure explains how to set the performance baseline expectations of a storage group to the characteristics currently measured for the previous two weeks.

Before you begin:

To perform this operation, a StorageAdmin role is required.

Setting new performance baseline expectations:

1. Select the storage system.
2. Do the following depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Select the storage group and click **View details** to open its details view.
2. If not already displaying, click the **SlO Compliance** tab.
3. Click **Reset Workload Plan**.
4. Review the **Current Scores** and the projected **New Baseline**.
5. If satisfied, click **OK**.

Once complete, the **Workload Planning** tab will update with the newly calculated performance metrics.

Analyzing FAST migration

The FAST Array Advisor wizard will guide you through the process of determining the performance impact of migrating the workload from one storage system (source) to another storage system (target). If the wizard determines that the target storage system can absorb the added workload, it will automatically create all the necessary auto-provisioning groups to duplicate the source workload on the target storage system.

Before you begin:

- To perform the analysis portion of this procedure, a Monitor role is required.
- To perform the migration preparation steps of this procedure, a StorageAdmin role is required.
- The following migrations are supported:

Source storage system	Target storage system
Enginuity 5874 - 5876	◆ Enginuity 5874 - 5876 or HYPERMAX OS 5977 or higher ◆ Target storage systems running Enginuity 5874 - 5876 must have at least one thin pool.
HYPERMAX OS 5977 or higher	HYPERMAX OS 5977 or higher

- The source storage system must be local and registered for Performance. Registering storage systems on page 630 explains how to registered storage systems.
- The storage group must:
  - Not be a child storage group. Only standalone or parent storage groups can be selected for analysis. If a parent storage group is selected, its child storage groups will be implicitly selected as well, and the analysis will apply to the entire collection of parent and child storage groups as a group.
  - Not have ongoing Workload Planning processing.
  - Be associated with a single masking view.
  - Only contain FBA volumes. It cannot be empty or contain only gatekeeper volumes.
  - Be associated with a Service Level Objective (HYPERMAX OS 5977), or associated with a FAST policy (Enginuity 5874 - 5876).

To analyzing FAST migration:

1. Select the storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.

   Note In eNAS operating environments, you can also perform this operation from **Storage Groups for File** page (**System › Dashboard › File Dashboard › File Storage Groups**).

   For Enginuity 5874 - 5876:

   Select **Storage › Storage Groups** to open the **Storage Group** list view.

1. Select the storage group, click >>, and select **FAST Array Advisor** to open the **FAST Array Advisor** wizard.
2. Select the **Target** storage system to include in the analysis, or select All available arrays to include all available storage systems in the analysis.
3. Do the following, depending on your target storage system selection:
   For target storage systems running HYPERMAX OS 5977 or higher:
   a. Select the Storage Resource Pool (**SRP**) to associate with the storage group.
   b. Select the Service Level Objective (**SLO**) to associate with the storage group. If the storage group is a parent, this involves selecting SLOs for each of the child storage groups.
   c. Click **Next**.
   d. Select the target storage system port configuration parameters to include in the analysis, including:
      - The **Number of ports** to use for the migrated storage group.
      - Whether to allow the system to select the most suitable ports (**All available ports**), or to select **Specific ports**. Selecting the latter will display list of ports from which you can select. The port list includes the port identifier and whether the VSA bit is enabled (Enabled is indicated with a flag).

   e. Click **Next**.

   For target storage systems running Enginuity 5874 - 5876:
   a. Select the FAST policy to associate with the storage group. If the storage group is a parent, this involves selecting policies for each of the child storage groups.
   b. Click **Next**.
   c. Select the target storage system port configuration parameters to include in the analysis, including:
      - The **Number of ports** to use for the migrated storage group.

- Whether to allow the system to select the most suitable ports (**All available ports**), or to select **Specific ports**. Selecting the latter will display list of ports from which you can select. The port list includes the port identifier and whether the VSA bit is enabled (Enabled is indicated with a flag).

d. Click **Next**.

For all available ports:

Click **Next**.

1. Determine the suitability of the migration:

   Note This option is only available under certain circumstances. For more information, refer to Suitability Check restrictions.

   Indicates a suitable migration.

   Indicates a non-suitable migration.

   In both cases, results are displayed in a bar chart by component (Front End, Back End, Cache) along with a score from 0 to 100 indicating the components expected availability on the target storage system after a potential migration.

   The current score for the component is shown in gray, with the additional load for the component shown in green or red indicating suitability. The additional score is red if the current and additional loads total more than 100.

   If the purposed migration is suitable, click **Next**; otherwise, click **Back** to configure another possible migration.

2. Prepare for the migration by reviewing the objects that will be created on the target storage system:

   Note Each object will be named exactly as it is on the source storage system. In the case of naming conflicts, a numeric suffix will be appended to target name to guarantee uniqueness. If the resulting suffixed name is longer than the 64 character limit, the target name will be truncated and the suffix added to the shortened name.

   - **Masking View** — Name of the masking view.
   - **Storage Group:**
     - **Storage Group** — Name of the storage group. If the storage group is a parent, you can expand the entry to view its child storage groups.
     - **Volumes** — Number of volumes to be created on the target storage system.
     - **Capacity (GB)** — Capacity of the storage group.

     Any host I/O limits set on the source storage group will be copied to the target storage group.
   - **Port Group:**

     Port Group — Name of the port group.

     Port — Name of the ports to be included in the port group.
   - **Initiator Group:**

- **Initiator Group** — Name of the initiator group. If the initiator group is a parent, you can expand the entry to view its child initiator groups.
- **Initiator** — Name of an initiator to be included in the initiator group.

The initiator group created on the target storage system will be an exact copy of the source workload's associated initiator group, including all initiators, child initiator group relationships, and flags. If this target cannot be created exactly as the source, then an error message will return.

3. Click **Finish**. This will create all the necessary auto-provisioning groups to duplicate the source workload on the target storage system. This may take some time to complete. The process has multiple steps, marked with a dialog box indicating the success or failure of each. Should the process fail, any partially created objects will remain on the target storage system until manually removed.

Volume management

Managing volumes (HYPERMAX OS 5977 or higher)

For storage systems running HYPERMAX OS 5977 or higher, the **Volumes** view provides you with a single place from which to view and manage all the volumes types on the system.

Note For instructions on managing volumes on storage systems running Enginuity versions 5671 - 5876, refer to Managing volumes (Enginuity 5671 - 5876) below.

To use the Volumes view:

1. Select the storage system.
2. Select **Storage › Volumes** to open the **Volumes** list view.

   By default, the view opens displaying thin volumes.
3. *Optional*: Filter the volume list by doing the following:
   a. Selecting/typing values for any number of the following criteria:
      - **Volume Configuration** — Filters the list for volumes with a specific configuration.
      - **Capacity equal to** — Filters the list for volumes with a specific capacity.
      - **Volume ID** — Filters the list for a volume with specific ID.
      - **Volume Identifier Name** — Filters the list for the specified volume name.
   b. Click **Find** to run the query now, or refine your query with conditional expressions, as described next.
   c. In the **Additional Criteria** section, select values for the following, and then click **Add Another**. Repeat this step for each additional expression.
      - **Category** — Specifies the broad category.
      - **Attribute** — Refines the category.
      - **Is equal to** — Establishes the relationship between the attribute and the value.
      - **Value** — Conditions for the specified attribute.
   d. Click **Find** to filter the volume list.

For field and control descriptions, refer to the following volume-specific help pages:

- **TDEV** — Viewing thin volumes
- **DATA** — Viewing DATA volumes on page 253

Managing volumes (Enginuity 5671 - 5876)

For storage systems running Enginuity versions 5671 - 5876, the **Volumes** dashboard provides a high-level view of all volumes on a storage system and a breakdown of the volume types.

Note For instructions on managing volumes on storage systems running HYPERMAX OS 5977, refer to Managing volumes (HYPERMAX OS 5977 or higher) on the previous page.



To use the Volumes view:

1. Select the storage system.
2. Select **Storage › Volumes** to open the **Volumes** dashboard.

   This dashboard contains two panels, All Volumes and Volumes Distribution Chart.
3. Expand the tree menu in the **All Volumes** panel to view the types/number of volumes on the Symmetrix system.
4. Double-click a volume type to open its list view, from which you can view details and perform operations on specific volumes.

All Volumes panel

The following properties display (Click a column heading to sort the list by that value):

- **General Volume Types** — High level volume container types such as, Regular, Virtual, Meta, Private,
- **Volume Type** — Name of volume type container.
- **Number of Volumes** — Number of volumes in the container.

The following controls are available:

- **Emulation** — Filters the listed volumes by emulation type (ALL, FBA, CKD). This filter only appears when there are multiple emulations types available.
- **View** — Allows for viewing the various volume container types.

Volumes Distribution Chart

The **Volumes Distribution Chart** provides a pie chart representation of volume type percentages of the total volumes and the count for each volume type.

Creating volumes

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. Do the following, depending on the storage operating environment and the type of volumes you are creating:
   HYPERMAX OS 5977 or higher:
   - TDEV — Creating thin volumes (HYPERMAX OS 5977) on page 217
   - Virtual Gatekeeper — Creating virtual gatekeeper volumes on page 219

     Note The maximum volume size supported on a VMAX3 system is 64 TB.

   Enginuity 5671 - 5876:
   - DATA — Virtual Provisioning on page 251
   - Diskless — Creating diskless volumes on the facing page
   - DRV — Creating DRV volumes on page 212
   - Gatekeeper — Creating gatekeeper volumes on page 213
   - Regular — Creating regular volumes on page 214
   - SAVE — Creating SAVE volumes on page 215
   - TDEV — Thin volumes on page 271
   - VDEV — Creating VDEV volumes on page 220

   In addition, you can also create volumes using a storage template, as described in Creating volumes using storage templates on page 219.

Creating DATA volumes

This procedure explains how to create DATA volumes on storage systems running Enginuity version 5671 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Private**.
4. Select **DATA** as the **Configuration**.
5. Select the **Disk Technology**.

   External disk technology is an option if the storage system has FTS (Federated Tiered Storage) enabled and available external storage.
6. Select the **Emulation** type.
7. Select the RAID **Protection** level.

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 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.
10. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
    - Click **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

The advanced options presented depend on the value selected for **Add to Pool**. Complete any of the following steps that are appropriate:

a. 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.

b. To enable the new volumes in the pool, select **Enable volume in pool**.

c. To rebalance allocated capacity across all the DATA volumes in the pool, select **Start Write Balancing**.

1. Click either of the following:
    - 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 851 and Previewing on page 850.
    - 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 on storage systems running Enginuity version 5671 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Virtual**.
4. Select **DLDEV** as the **Configuration**.
5. Select the **Emulation** type.
6. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter the volume capacity.

7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **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. If Auto Meta is enabled on the system, and if you are attempting to create volumes larger than the **Minimum Meta Capacity**, specify values for the following in the **Define Meta** panel:
   - **Member capacity (Cyl/MB/GB)** — Size of the meta members to use when creating the meta volumes.
   - **Configuration (Striped/Concatenated)** — Whether to create striped or concatenated meta volumes.
3. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating DRV volumes

This procedure explains how to create DRV volumes on Symmetrix systems running Enginuity version 5773 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Private**.
4. From **Configuration**, select **DRV**.
5. Select the **Disk Technology** on which to create the volumes.
6. Select the **Emulation** type.
7. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.

8. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **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. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating gatekeeper volumes

This procedure explains how to create gatekeeper volumes on storage systems running Enginuity version 5773 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. To create a gatekeeper, complete the following steps:
   a. In **Volume Type**, click **Regular**.
   b. Select **Gatekeeper** as the **Configuration**.

   To create a virtual gatekeeper, complete the following steps:
   a. In **Volume Type**, click **Virtual**.
   b. Select **Virtual Gatekeeper** as the **Configuration**.
4. Select the **Emulation** type.
5. Type the **Number of Volumes** to create.
6. 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).
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating regular volumes

This procedure explains how to create regular volumes on storage systems running Enginuity version 5671 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In the Volume Type panel, click **Regular**.
4. Select the **Configuration**.
5. Select the **Disk Technology**.

   External disk technology is an option if the storage system has FTS (Federated Tiered Storage) enabled and available external storage.
6. Select the **Emulation** type.
7. Select the RAID **Protection** level.
8. Specify the capacity to create by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
9. Click one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. z/OS Only: Type the **SSID** for the new volume, or click **Select...** to open a dialog from which you can select an SSID.This is required for volumes on storage systems with ESCON or FICON directors (or mixed systems). This field is optional on storage systems running Enginuity 5773 or higher when reducing the number of mirrors.
2. To create the volumes from a specific **Disk Group**, select one (disk group number and name).
3. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name:**
   - **None** — Allows the system to name the volumes (Default).
   - **Name Only** — All volumes will have the same name.
   - **Name + VolumeID** — All volumes will have the same name with a unique storage system 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 231

4. 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.
5. If Auto Meta is enabled on the system, and if you are attempting to create volumes larger than the **Minimum Meta Capacity**, specify values for the following in the **Define Meta** panel:
   - **Member capacity (Cyl/MB/GB)** — Size of the meta members to use when creating the meta volumes.
   - **Configuration (Striped/Concatenated)** — Whether to create striped or concatenated meta volumes.
6. Click either of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating SAVE volumes

This procedure explains how to create SAVE volumes on Symmetrix systems running Enginuity version 5773 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Private**.
4. 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.
5. Select the **Disk Technology**.

   External disk technology is an option if the Symmetrix system has FTS (Federated Tiered Storage) enabled and available external storage.
6. Select the **Emulation** type.
7. Select the RAID **Protection** level.
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 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.

10. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
    - Click **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. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating thin volumes (Enginuity 5773 - 5876)

This procedure explains how to create thin volumes on storage systems running Enginuity version 5773 - 5876. For instructions on creating thin volumes on storage systems running HYPERMAX OS 5977 or higher, refer to Creating thin volumes (HYPERMAX OS 5977) on the facing page.

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Virtual**.
4. Select **Configuration (TDEV or BCV + TDEV)** or thin volumes.
5. Select the **Emulation** type.
6. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
7. 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).
8. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **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:**
   - **None** — Allows the system to name the volumes (Default).
   - **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 231
2. To **Allocate Full Volume Capacity**, select the option.
3. If you selected to allocate capacity in the previous step, you can mark the allocation 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.
4. 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.
5. If Auto Meta is enabled on the system, and if you are attempting to create volumes larger than the **Minimum Meta Capacity**, specify values for the following in the **Define Meta** panel:
   - **Member capacity (Cyl/MB/GB)** — Size of the meta members to use when creating the meta volumes.
   - **Configuration (Striped/Concatenated)** — Whether to create striped or concatenated meta volumes.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating thin volumes (HYPERMAX OS 5977)

This procedure explains how to create thin volumes on VMAX systems running HYPERMAX OS 5977. For instructions on creating thin volumes on storage systems running Enginuity 5773 - 5876, refer to Thin volumes on page 271.

1. Select the storage system.

2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. Select TDEV as the **Configuration**.
4. Select the **Emulation** type.
5. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
6. *Optional*: To add the volumes to a storage group, click **Select**, select the volumes, and then click **OK**.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. If creating thin volumes or a thin BCVs, you can specify to **Allocate Full Volume Capacity**. In addition, 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.
2. To name the new volumes, select one of the following **Volume Identifiers** and type a **Name:**
   - **None** — Allows the system to name the volumes (Default).
   - **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 231
3. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating virtual gatekeeper volumes

Before you begin:

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

Creating virtual gatekeepers:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Virtual**.
4. Select **Virtual Gatekeeper** as the **Configuration**.
5. *Optional*: Select the **Emulation** type.
6. Type the **Number of Volumes**.
7. *Optional*: To add the volumes to a storage group, click **Select**, select the storage group, and then click **OK**.
8. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating volumes using storage templates

This procedure explains how to use storage templates to create volumes on Symmetrix systems running Enginuity 5773 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Template**.
4. 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.
5. 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.
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 851 and Previewing on page 850.
   - 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:**
   - **None** — Allows the system to name the volumes (Default).
   - **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 231
2. Click either of the following:
   - **Add to Job List**. to create the volumes at a later time, as described in Scheduling jobs (page 851), or from the job list as described in Viewing the job list (page 853).
   - Expand **Add to Job List** and click **Run Now** to create the volumes now.

Creating VDEV volumes

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Virtual**.
4. Select **VDEV** as the **Configuration**.

5. Select the **Emulation** type.
6. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**.

   If Auto meta is enabled on the system then it displays as enabled with a green check mark.
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 851 and Previewing on page 850.
   - 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. 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.
2. If Auto Meta is enabled for the system, and if you are attempting to create volumes larger than the **Minimum Meta Capacity**, specify values for the following in the **Define Meta** panel:
   - **Member capacity (Cyl/MB/GB)** — Size of the meta members to use when creating the meta volumes.
   - **Configuration (Striped/Concatenated)** — Whether to create striped or concatenated meta volumes.
3. Click one of the following:
   - **Add to Job List** to create the volumes at a later time, as described in Scheduling jobs on page 851, or from the job list as described in Viewing the job list on page 853.
   - Expand **Add to Job List** and click **Run Now** to create the volumes now.

Deleting volumes

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select the volumes and click **Delete**.
5. Click **Delete** in the confirmation dialog box.

Duplicating volumes

The following explains how to duplicate volumes.

Before you begin:

- You cannot duplicate RDF, SFS, or VAULT volumes.
- 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 storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select the volume, click more >> , and click **Duplicate Volume** 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 851), or from the job list as described in Viewing the job list (page 853).
   - 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:**
   - **None** — Allows the system to name the volumes (Default).
   - **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 231

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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Assigning Symmetrix priority

This procedure explains how to prioritize the service time of the host I/O to individual volumes or groups of volumes (DGs or SGs).

Before you begin:

This feature requires Enginuity 5671 - 5876.

To assign host priority to individual volumes:

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

To assign host priority to groups of volumes:

1. Select the storage system.
2. To assign priority to storage groups, select **Storage › Storage Groups** to open the **Storage Groups** list view.

   To assign priority to device groups, select **Data Protection › Replication Groups and Pools › Device Groups** to open the **Device Group** list view.
3. Select the group, 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 storage systems running Enginuity 5874 or higher, you cannot increase or decrease the mirror protection of a volume.
- When adding DRV attributes, volumes must be unmapped.
- Full swap operations require the R1 and R2 devices to be the same size.
- Only the head of a metavolume can have its type changed. The metamembers will automatically have the changes applied.
- You cannot convert one member of a raidset to unprotected without converting all the members to unprotected.
- When adding/removing SRDF attributes, there are no restrictions on I/O. The SRDF pair must be split or failed over. If failed over, the R1 device must be unmapped.
- When adding/removing BCV attributes, there are no restrictions on I/O. The standard/BCV pair must be split.

To change the configuration of volumes:

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
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 Enginuity 5773 or higher when reducing the number of mirrors.
8. To reserve the volumes, select **Reserve Volumes**. 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. Do one of the following:
   - Click **Add to Job List** to schedule the job for a later time, as described in Scheduling jobs (page 851).
   - Expand **Add to Job List** and click **Run Now** run the job now.

Expanding existing volumes

This procedure explains how to expand the capacity of existing, externally-visible thin volumes.

Before you begin:

- Volume expansion can be performed only for HYPERMAX OS 5977, or higher.
- Volumes can be expanded up to, and including, a size of 64 TB.
- Volume expansion cannot be performed if any of the following operations are currently being carried out on the volume:
  - free all
  - reclaim
  - allocation
- Restriction apply if the volume
  - is a gatekeeper
  - is an ACLX volume
  - has Celerra FBA emulation
  - is an AS400 volume
  - is FTS encapsulated
  - has a SnapVX session defined, or is being replicated

To set volume status for individual volumes:

1. Select the storage system.
2. To open the Expand Volume dialog box, do one of the following:

*Open the **Expand Volume** dialog box from the **Storage Group Volumes** list*

a. Click **Storage › Storage Groups Dashboard**.

b. Click **Total**.

c. Select the storage group and click **View Details**.

d. In the **Related Objects** pane, click **Volumes**. The **Storage Group Volumes List** view displays.

e. Do one of the following:

- Select a volume and click **Expand Volume**.
- Select a volume and click **View Details** and then click >> **› Expand Volume**.

*Open the **Expand Volume** dialog box from the **Volumes** list*

a. Click **Storage › Volumes**.

b. Select a volume and click >> **› Expand Volume**.

The **Expand Volume** dialog box displays.

1. In the **Volume Capacity** field of the **Expand Volume** dialog box, type or select the new capacity of the volume. The **Total Capacity** and **Additional Capacity** figures update automatically.

2. To reserve the volume, select **Reserve Volumes**.
3. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Mapping volumes

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes and click >> and **Map** to open the mapping wizard.
5. Select one or more **Ports**.

Note When performing this operation on storage systems running HYPERMAX OS 5977 or higher, only ACLX-disabled ports will be available for selection.

6. To reserve the volumes, select **Reserve Volumes**. In addition you can also type reserve **Comments** and select an **Expiration**. The default values for **Reserve Volumes**and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
7. Click **Next**.
8. To change an automatically generated LUN address, do the following; otherwise, click **Next** to accept the generated address.
   a. Double-click the address to open the **Set Dynamic LUN Address** dialog box.
   b. To use a new **Starting LUN**, double-click it and type a new address over it, or select an address and click **Next Available LUN** to increment the generated address to the next available address. When done, click **Apply Starting LUN**.
   c. Click **OK** to return to the mapping wizard
   d. Click **Next**.
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. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Unmapping volumes

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes and click >> and **Unmap** to open the **Unmap Volumes** wizard.
5. Select one or more ports.
6. To reserve the volumes, select **Reserve Volumes**. In addition you can also type reserve **Comments** and select an **Expiration**. The default values for **Reserve Volumes** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
7. Click **Next**.
8. 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.
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 851 and Previewing on page 850.
   - 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 storage systems running Enginuity 5876.163.105 or higher. This feature is not supported on storage systems running HYPERMAX OS 5977 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. However, starting with Enginuity 5876.280, you can use optimized read miss without a XtremeSW Cache Adapter. To use optimized read miss without the adapter, you must set the **Optimized Read Miss** mode to **On**.

To set optimized read miss:

1. Select the storage system.
2. Select **Storage › Volumes** to open the **Volumes** dashboard.
3. In the **All Volumes** panel, double-click the volume type to open its list view.
4. Select one or more volumes, click more >> , and select **Set Optimized Read Miss** to open the **Optimized Read Miss** dialog box.
5. *Optional*: Click **Show Selected Volumes** to view details on the selected volumes.
6. Select a **Set Optimized Read Miss** mode:
   - **System Default** — Storage system determines whether to enable or disable optimized read miss mode for the specified volumes/group.
   - **Off** — Disables optimized read miss mode, regardless of the configuraion.
   - **On** — Enables optimized read miss mode for both XtremCache and non-XtremCache EFD-only configurations.
7. 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 storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >> , and select **Set Volume Status** to open the **Set Volume Status** dialog box.
5. To view details on the selected volumes, click **Show Current Volumes**.
6. **Set Volume 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.
   - **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.
7. *Optional:* For HYPERMAX OS 5977 or higher, select **SRDF/Metro**.
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 storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >> , and select **Set Volume Attributes** to open the **Set Volume Attributes** dialog box.
5. To view details on the selected volumes, click **Show Selected Volumes**.
6. 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. This feature only applies/appears for Symmetrix systems running Enginuity 5773-5876. 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.
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. Do one of the following:
   - 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 851 and Previewing on page 850.

- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Setting volume identifiers

1. Select the storage system.
2. Do one of the following:

*Open the **Set Volume Identifiers** dialog box from the **Volumes** list*

a. Select **Storage › Volumes**.

1. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.

b. Select one or more volumes, click more >> , and select **Set Volume Identifiers**.

*Open the **Set Volume Identifiers** dialog box from the **Protocol Endpoints** list*

a. Select **Storage › VVols Dashboard**.

b. Click **Protocol Endpoints**. The **Protocol Endpoints** list displays.

c. Select one or more protocol endpoints, and click **Set Volume Identifiers**.

The **Set Volume Identifiers** dialog box displays.

1. To view details on the selected volumes, click **Show Current Volumes**.
2. 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.
3. 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.
4. Type the **Volume 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.
5. To reserve the volumes, select **Reserve Volumes**. In addition you can also type reserve **Comments** and select an **Expiration** date. The default values for **Reserve Volumes** and **Comments** are set in Symmetrix Preferences for volumes reservations. If the volumes are not automatically reserved you can optionally reserve them here.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

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.

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 storage volumes.

To set the copy pace at the storage group level:

1. Select the storage system.
2. Do the following, depending on the storage operating environment:
   For HYPERMAX OS 5977 or higher:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **SLO Compliance** panel, click **Total** to open the **Storage Groups** list view.
   For Enginuity 5874 - 5876:
   Select **Storage › Storage Groups** to open the **Storage Group** list view.
1. Perform one of the following actions:
   - *All volumes in the storage group*: Select the storage group, click more >> , and select **Replication QoS** to open the **Set Replication Priority QoS** dialog box.
   - *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**.
2. Select the **Operation Type** from the possible values:
   - **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. This option is not available on storage systems running HYPERMAX OS 5977 or higher.

3. 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 storage 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 storage system is running an Enginuity version earlier than 5875.
4. Click **OK**.

To set the copy pace at the device group level:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › Device Groups** to open the **Device Groups** view.
3. Select the device group, click more >> , and select **Replication QoS** to open the **Set Replication Priority QoS** dialog box.
4. Select the **Operation Type** from the possible values:
   - **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. This option is not available on storage systems running HYPERMAX OS 5977 or higher.
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 storage 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.
6. *If performing this operation on a group*: Select the type of volumes on which to perform the operation.
7. Click **OK**.

To set the copy pace at the volume level:

1. Select the storage system.
2. Select **Storage › Volumes** to open the **Volumes** dashboard.
3. In the **All Volumes** panel, double-click the volume type to open its list view.

4. Select one or more volumes, click more >> , and select **Replication QoS** to open the **Replication QoS** dialog box.
5. Select the **Operation Type** from the possible values:
   - **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. This option is not available on storage systems running HYPERMAX OS 5977 or higher.
6. 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 storage 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.
7. Click **OK**.

Managing Meta Volumes

Creating meta volumes

Before you begin:

- Meta volumes are supported on storage systems running Enginuity 5773 - 5876.
- On storage systems running Enginuity versions lower than 5875, only unmapped thin volumes (hypers) can be formed into meta volumes.
- On storage 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 *Solutions Enabler Array Controls CLI Guide*.
- When creating meta volumes, Unisphere for VMAX will attempt to instill best practices in the creation process by setting the following defaults in the **Create Meta Volume** wizard:
  - **Meta Volume Configuration** = Striped
  - **Meta Volume Member Count including Head** = 8

  Note that these best practices do not apply to volumes created with the CKD-3390 emulation type.

To create meta volumes:

1. Select the storage 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 Meta Volume** to open the **Create Meta** wizard.
5. Select the **Emulation** type.
6. If creating FBA volumes, select whether to create them from **Create Volumes**or **Use Existing Volumes** 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 (Request/Available)** containing the meta volumes.

    g. To reserve the volumes, select **Reserve Volumes**.

    h. Click **Next**.

    By Manually Selecting Existing Volumes (Advanced):

    a. Select from the listed volumes.

    b. To reserve the volumes, select **Reserve Volumes**.

    c. Click **Next**.

    Using New Standard Provisioned Volumes:

    a. Specify the **Number of Meta Volumes**.

    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 a **Disk Group**.

g. If you are creating CKD meta volumes, type or select an **SSID**.

h. Click **Next**

Using New Virtual Volumes:

a. Specify the **Number of Meta Volumes**.

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. 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 **Run Now**.

Adding meta members

Before you begin:

- Meta volumes are supported on storage systems running Enginuity 5773 - 5876.
- To expand a bound striped thin meta volume on a storage system running an Enginuity version lower than 5875, you must first unbind the volume from the pool. This operation is allowed on storage systems running Enginuity 5875 or higher without having to unbind the volume, however, you must select the **Protect Data** option.
- When expanding meta thin volumes with BCV meta protection, the volumes must be fully allocated to a pool and they must have the **Persist preallocated capacity through reclaim or copy** option set on them. This is because binding thin meta BCV volumes is done through the pool and not through the thin BCV volume selection. For more information on allocating thin pool capacity for thin volumes, refer to Managing thin pool allocations on page 263.

To add meta members:

1. Select the storage 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. 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 **Add to Job List** or **Run Now** and refer to Managing jobs on page 848.

Removing meta members

Before you begin:

- Meta volumes are supported on storage systems running Enginuity 5773 - 5876.
- You can only remove members from concatenated meta volumes.

To remove meta members

1. Select the storage 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 **Add to Job List** or **Run Now**.

Converting meta volumes

This procedure explains how to change the configuration of a meta volume.

Before you begin:

Meta volumes are supported on storage systems running Enginuity 5773 - 5876.

To convert meta volumes:

1. Select the storage 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. To reserve the volumes, select **Reserve 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Dissolving meta volumes

Before you begin:

Meta volumes are supported on storage systems running Enginuity 5773 - 5876.

To dissolve meta volumes:

1. Select the storage 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. *Optional*: If required, select **Reserve Volumes** and/or **Delete meta members after dissolve**. Note that selecting **Delete meta members after dissolve** requires the operation to be run immediately and not scheduled.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing meta volumes

Before you begin:

Meta volumes are supported on storage systems running Enginuity 5773 - 5876.

To view meta volumes:

1. Select the storage 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** — Volume number.
- **Type** — Volume type.
- **Meta Config** — Volume configuration.
- **Stripe 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 Meta Volume** — See Managing Meta Volumes on page 235.
- **Add Member** — See Adding meta members on page 238.
- **Dissolve** — See Dissolving meta volumes on the previous page.
- **Convert** — See Converting meta volumes on page 239.
- **View Details** — See Viewing meta volume details on the next page.
- **Tag for RecoverPoint** — See Tagging and untagging volumes for RecoverPoint on page 518.
- **Untag for RecoverPoint** — See Tagging and untagging volumes for RecoverPoint on page 518.
- **Map** — See Mapping volumes on page 226.
- **Unmap** — See Unmapping volumes on page 227.
- **Change Volume Configuration** — See Changing volume configuration on page 223.
- **Duplicate Volume** — See Duplicating volumes on page 221.
- **Set Volume Attributes** — See Setting volume attributes on page 230.
- **Set Volume Identifiers** — See Setting volume identifiers on page 231.
- **Set Volume Status** — See Setting volume status on page 229.
- **Replication QOS** — See Setting replication QoS on page 232.
- **Assign Dynamic Cache Partition** — See Assigning dynamic cache partitions on page 875.
- **Pin** — See Pinning and unpinning volumes on page 197.
- **Unpin** — See Pinning and unpinning volumes on page 197.

- **VLUN Migration** — See Migrating regular volumes on page 485.
- **Set Optimized Read Miss** — See Setting optimized read miss on page 227.

Viewing meta volume details

Before you begin:

Meta volumes are supported on storage systems running Enginuity 5773 - 5876.

To view meta volume details:

1. Select the storage 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.
- **Physical Name** — Physical name.
- **Volume Identifier** — Volume identifier.
- **Type** — Volume configuration.
- **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.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **Capacity (GB)** —Volume capacity in GBs.
- **Capacity (MB)** —Volume capacity in MBs.
- **Capacity (Cylinder)** —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
- **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 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.
- **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.
- **XtremSW Cache Attached** — Indicates whether XtremSW cache is attached to the volume.
- **Optimized Read Miss** — Cacheless read miss status.

The following controls are available:

- **Create Meta Volume** — See Managing Meta Volumes on page 235.
- **Add Member** — See Adding meta members on page 238.
- **Dissolve** — See Dissolving meta volumes on page 240.
- **Convert** — See Converting meta volumes on page 239.
- **Tag for RecoverPoint** — See Tagging and untagging volumes for RecoverPoint on page 518.
- **Untag for RecoverPoint** — See Tagging and untagging volumes for RecoverPoint on page 518.
- **Map** — See Mapping volumes on page 226.
- **Unmap** — See Unmapping volumes on page 227.
- **Change Volume Configuration** — See Changing volume configuration on page 223.
- **Duplicate Volume** — See Duplicating volumes on page 221.
- **Set Volume Attributes** — See Setting volume attributes on page 230.
- **Set Volume Identifiers** — See Setting volume identifiers on page 231.
- **Set Volume Status** — See Setting volume status on page 229.
- **Assign Dynamic Cache Partition** — See Assigning dynamic cache partitions on page 875.
- **Replication QOS** — See Setting replication QoS on page 232.
- **Pin** — See Pinning and unpinning volumes on page 197.
- **Unpin** — See Pinning and unpinning volumes on page 197.
- **VLUN Migration** — See Migrating regular volumes on page 485.
- **Set Optimized Read Miss** — See Setting optimized read miss on page 227.

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 storage 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 storage 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** view.
5. Select a disk from the list and click **Remove Disk**.
6. Click **OK** 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 storage system.
2. Select **Storage › Disk Groups** to open the **Disk Groups** list view.
3. Select one or more disk groups and click **Delete Disk Group**.
4. Click **Delete** in the confirmation message.

Viewing disk groups

1. Select the storage 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.
   - **Disk Location** — Indicates whether disk is internal or external.
   - **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.

   The following controls are available:

   - **View Details** — See Viewing disk group details on the facing page.
   - **Delete Disk Group** — See Deleting disk groups on the previous page.

Viewing disk group details

1. Select the storage 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.
- **Speed (RPM)** — Speed of the disks in the group.
- **Form Factor** — Physical size of the disks in the group.
- **Disk Location** — Whether the disks in the group are internal or external.

The following controls are available:

- **Rename** — See Renaming disk groups on page 245.
- **Delete** — See Deleting disk groups on page 245.

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 storage 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 storage 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** list view.

   Use the **Disks** 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 Capacity (GB)** — Disk capacity.
   - **Used (%)** — Percent of disk capacity.

   The following controls are available:

   - **View Details** — See Viewing disk details below.
   - **Remove Disks** — See Removing disks from disk groups on page 245.

Viewing disk details

1. Select the storage 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 ID** — Product 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.
- **Total Capacity (GB)** — Useable disk capacity in Gigabytes.
- **Free Capacity (GB)** — Free disk capacity in Gigabytes.
- **Actual Capacity (GB)** — Actual disk capacity in Gigabytes.
- **Used (%)** — Percentage of used disk capacity to the total disk capacity.
- **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 245.

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 on storage systems running Enginuity version 5671 - 5876:

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Private**.
4. Select **DATA** as the **Configuration**.
5. Select the **Disk Technology**.

   External disk technology is an option if the storage system has FTS (Federated Tiered Storage) enabled and available external storage.
6. Select the **Emulation** type.
7. Select the RAID **Protection** level.
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 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.
10. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
    - Click **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

The advanced options presented depend on the value selected for **Add to Pool**. Complete any of the following steps that are appropriate:

a. 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.

b. To enable the new volumes in the pool, select **Enable volume in pool**.

c. To rebalance allocated capacity across all the DATA volumes in the pool, select **Start Write Balancing**.

1. Click either of the following:
   - 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 851 and

Previewing on page 850.

- Expand **Add to Job List**, and click **Run Now** to perform the operation 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 storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
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 storage 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 storage 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 storage 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 storage system.
2. Select **Storage › Volumes** to open the **Volumes** view.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.

   The **DATA** volumes list view allows you to view and manage DATA volumes.

The following properties display:

Note Depending on the storage environment, some of the following properties may not apply.

- **Name** — Volume name.
- **Type** — Volume configuration.
- **Pool Name** — Pool in which the volume resides.
- **Pool Type** — Type of pool in which the volume resides.
- **Capacity (GB)**—Volume capacity.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **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:

Note Depending on the storage environment, some of the following controls may not be available.

- **Create Volume** — See Virtual Provisioning on page 251.
- **Delete** — See Deleting volumes on page 221.
- **Enable** — See Enabling and disabling DATA volumes on page 252.
- **Disable** — See Enabling and disabling DATA volumes on page 252.
- **Activate** — See Activating and deactivating DATA volumes on page 252.
- **Deactivate** — See Activating and deactivating DATA volumes on page 252.
- **View Details** — See Viewing DATA volume details below.
- **Start Draining** — See Start draining DATA volumes on the previous page.
- **Stop Draining** — See Stop draining DATA volumes on the previous page.
- **Assign Dynamic Cache Partition** — Assigning dynamic cache partitions on page 875
- **Duplicate Volume** — Duplicating volumes on page 221

Viewing DATA volume details

1. Select the storage system.
2. Select **Storage › Volumes**

3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. 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:

Note Depending on the storage environment, some of the following properties may not apply.

- **Name** — Volume name.
- **Physical Name** — Physical name.
- **Volume Identifier** — Volume identifier.
- **Type** — Volume configuration.
- **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.
- **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.
- **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 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.
- **RecoverPoint tagged** — Indicates whether 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 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:

Note Depending on the storage environment, some of the following controls may not be available.

- **Create Volume** — Virtual Provisioning (page 251)
- **Delete** — Deleting volumes (page 221)
- **Enable** — Enabling and disabling DATA volumes on page 252
- **Disable** — Enabling and disabling DATA volumes on page 252
- **Activate** — See Activating and deactivating DATA volumes on page 252.
- **Deactivate** — See Activating and deactivating DATA volumes on page 252.
- **Start Draining** — See Start draining DATA volumes on page 253.
- **Stop Draining** — See Stop draining DATA volumes on page 253.
- **Assign Dynamic Cache Partition** — Assigning dynamic cache partitions on page 875
- **Duplicate Volume** — Duplicating volumes on page 221

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.
- When creating thin pools, Unisphere for VMAX will attempt to instill best practices in the creation process by updating the default **Protection** level according to the selected **Disk Technology**:

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

To create a thin pool:

1. Select the storage system.

2. Select **Storage › Thin Pools** to open the **Thin Pools** list view.
3. Click **Create Thin Pool** 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 (0-65534)** and type a percentage.
3. Type the **Rebalancing Variance (1-50)**. 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 (2-1024)**. This is the maximum number of volumes in the pool on which the rebalancing algorithm will concurrently operate. Possible values range from 2 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 (1-80)** 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 80.
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 851, or from the job list as described in Viewing the job list on page 853.
   - Expand **Add to Job List** and click **Run Now** to create the pool now.

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.

Before you begin:

Unisphere for VMAX supports best practices, which state that volumes from different drive technologies should not be mixed in the same thin pool. To this end, Unisphere for VMAX will only expand a thin pool with volumes from the same disk group as the volumes already in the pool. This is an important distinction from Solutions Enabler, which does not impose this restriction.

To expand a thin pool:

1. Select the storage system.
2. Select **Storage › Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **Expand** 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** and select the unit of capacity.
   - **Percentage** — Type an amount in the **Percentage Increase** field.
5. *Optional*: Click **Show Advanced** and complete the following steps:
   - Select **Volume Capacity**, in GB.
   - Click **Start Write Balancing** 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**.

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 252.
- The drain must not cause the enabled volumes to end up with greater than 90% utilization in the pool. To calculate this, 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 ofthe 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 storage 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** 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 *VMAX 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 Enginuity 5874.207.166 or higher.

To write balance a thin pool:

1. Select the storage 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**.

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 storage 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.

Adding/Removing thin pool members

This procedure explains how to add/remove members from a thin pool.

Before you begin:

- The Symmetrix system must be running Enginuity 5773 - 5876.
- Before you can remove a thin pool member (DATA volume), you must first disable it.
- Unisphere for VMAX supports best practices, which state that volumes from different drive technologies should not be mixed in the same thin pool. To this end, the Add Volumes to Thin Pool dialog box will only allow you to add volumes from the same disk group as the volumes already in the pool. This is an important distinction from Solutions Enabler, which does not impose this restriction.

To add/remove thin pool members:

1. Select the storage 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** 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** wizard.

   b. Locate the volumes by selecting/typing values for any number of the following criteria:

      - **Capacity equal to** — Filters the list for volumes with a specific capacity.
      - **Volume ID** — Filters the list for a volume with specific ID.
      - **Volume Identifier Name** — Filters the list for the specified volume name.

   c. Click **Find Volumes** to run the query now, or specify **Additional Criteria** to add conditional expressions to your query, as described next.

      1. Select values for the following:
         - **Category** — Specifies the broad category.
         - **Attribute** — Refines the category.
         - **Operator** — Establishes the relationship between the attribute and the value.
         - **Value** — Conditions for the specified attribute.
      2. Click **Add Another** and repeat the previous step for each additional expression.
      3. Click **Find Volumes**.

   d. In the **Available Volumes** table, select the volumes.

   e. *(Optional):* Select **Start Write Balancing**.

   f. Click **Add to Pool**.

   Removing volumes:

   a. Select one or more volumes and click **Remove**.

   b. Click **OK** in the confirmation message.

Enabling and disabling thin pool members

1. Select the storage 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.

Managing thin pool allocations

The following describes how to start and stop allocating thin pool capacity from the **Volumes** view. In addition, you can also perform this operation from the following views:

- Storage Groups (HYPERMAX OS 5977 or higher): **Storage › Storage Groups Dashboard › Storage Groups**

- Storage Groups (Enginuity 5874 - 5876): **Storage › Storage Groups**
- Device Groups: **Data Protection › Replication Groups and Pools › Device Groups**
- File Storage Groups: **System › Dashboard › File Dashboard › File Storage Groups**

Before you begin:

- You can only allocate thin pool capacity to bound thin volumes.
- Thin pool allocation for individual thin volumes requires Enginuity 5773 - 5876, or HYPERMAX OS 5977 or higher.
- Thin pool allocation for thin volumes in a storage group requires Enginuity 5875 - 5876, or HYPERMAX OS 5977 or higher.
- This procedure explains how to perform this operation from the **Volumes** view. You can also perform this procedure from storage group views. Depending from where you are performing this procedure, some of the following steps may not apply.

To start thin pool allocation:

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >> , and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
5. Select **Allocate Volumes**. In addition, on storage systems running Enginuity 5875.139.93 - 5876 or HYPERMAX OS 5977 or higher, 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.
6. 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.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

To stop thin pool allocation:

1. Select the storage system.

2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >>, and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
5. Select **Stop Allocate**.
6. 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.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

For more information about thin pools and thin provisioning concepts, refer to the *VMAX Product Guide*.

Managing space reclamation

The following describes how to start and stop the space reclamation option, freeing unused pool capacity and reclaiming unwritten tracks from thins volumes. In addition, you can also perform this operation from the following views:

- Storage Groups (HYPERMAX OS 5977 or higher): **Storage › Storage Groups Dashboard › Storage Groups**
- Storage Groups (Enginuity 5874 - 5876): **Storage › Storage Groups**
- Device Groups: **Data Protection › Replication Groups and Pools › Device Groups**
- File Storage Groups: **System › Dashboard › File Dashboard › File Storage Groups**

Before you begin:

- You can only reclaim thin pool capacity from bound thin volumes.
- Thin pool reclamation for individual thin volumes requires Enginuity 5874 - 5876 or HYPERMAX OS 5977 or higher.
- Thin pool reclamation for thin volumes in a storage group requires Enginuity 5875 or higher.
- This procedure explains how to perform this operation from the **Volumes** view. You can also perform this operation from storage group views. Depending from where you are performing this procedure, some of the following steps may not apply.

To start space reclamation:

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >>, and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
5. To start reclaiming unwritten tracks and those written completely with zeros, select **Reclaim Volumes**. This option is only available on storage systems running Enginuity 5875.139.93 or higher, or storage systems running HYPERMAX OS 5977 or higher.
6. *Optional*: To reclaim tracks that are unwritten or zero-based, even if they are marked as persistent, select **Reclaim persistent capacity**.
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. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

To stop space reclamation:

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >>, and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
5. Select one or more volumes, click more >> and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
6. Select **Stop Reclaim**. In addition, on storage 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.
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. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

For more information about thin pools and thin provisioning concepts, refer to the *VMAX Product Guide*.

Managing thin pool capacity

The following describes how to start and stop the process of freeing allocated thin pool capacity from the **Volumes** view. In addition, you can also perform this operation from the following views:

- Storage Groups (HYPERMAX OS 5977 or higher): **Storage › Storage Groups Dashboard › Storage Groups**
- Storage Groups (Enginuity 5874 - 5876): **Storage › Storage Groups**
- Device Groups: **Data Protection › Replication Groups and Pools › Device Groups**
- File Storage Groups: **System › Dashboard › File Dashboard › File Storage Groups**

Before you begin:

- You can only reclaim thin pool capacity from bound thin volumes.
- Thin pool reclamation for individual thin volumes requires Enginuity 5874 - 5876 or HYPERMAX OS 5977 or higher.
- This procedure explains how to perform this operation from the **Volumes** view. You can also perform this operation from storage group views. Depending from where you are performing this procedure, some of the following steps may not apply.

To start freeing unused capacity:

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >> , and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
5. Select **Free Volumes**.
6. *Optional*: To free all allocations associated with the volumes, regardless of whether the data is written, select **Free all allocations (written and unwritten)**. This option is only available on storage systems running HYPERMAX OS 5977 or higher.

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. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

To stop freeing unused capacity:

1. Select the VMAX system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select one or more volumes, click more >> , and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
5. Select **Free Volumes**. In addition, on storage 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 storage systems running Enginuity 5875.139.93.
6. 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.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

For more information about thin pools and thin provisioning concepts, refer to the *Solutions Enabler Symmetrix Array Management CLI Product Guide*.

Viewing thin pools

1. Select the storage 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.
   - **Configuration** — Configuration of the pool.
   - **Emulation** — Emulation of the pool.
   - **Allocated Capacity** — Percentage of the pool that is allocated.
   - **Enabled Capacity (GB)**— Capacity of the pool in GB.

   The following controls are available:

   - **Create Thin Pool** — See Thin pools on page 257.
   - **Expand** — See Expanding thin pools on page 259.
   - **View Details** — See Viewing thin pool details below.
   - **Delete** — See Deleting thin pools on page 261.

Viewing thin pool details

1. Select the storage 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.
- **Type** — The pool type.
- **Technology** — Disk technology on which the pool resides.
- **Disk Location** — Whether the disk group is Internal to the storage system or an External storage system or storage device.

- **Emulation** — Emulation type for the pool.
- **Total Capacity (GB)** — Total capacity of the pool.
- **Free Capacity (GB)** — 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.
- **Set Maximum Subscription** — Enable oversubscription for the pool.
- **Maximum Subscription (0-65534)** — Acceptable oversubscription ratio for the pool.
- **% Subscription** — Current subscription percentage.
- **Enabled Capacity (GB)** — Sum of capacity of all enabled DATA volumes in the pool.
- **Allocated Capacity (GB)** — Pool capacity allocated to thin volumes.
- **State** — Pool state (Enabled, Disable, Balancing).
- **Rebalance Variance (0-50)**— 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 Reserved Capacity Enabled** — Whether a percentage of the capacity of the thin pool is reserved.
- **Pool Reserved Capacity** — The percentage of the capacity of the thin pool that will be reserved for non-FAST activities
- **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.
- **Saved By Compression (Tracks)** — The number of tracks saved by using VP Pool Compression.

- **Pool Compression Ratio (%)** — Percentage of allocated pool capacity that is compressed.

The following controls are available:

- **Expand** — See Expanding thin pools on page 259.
- **Bind** — See Binding/Unbinding/Rebinding thin volumes on page 273.
- **Stop Write Balancing** — See Starting and stopping thin pool write balancing on page 261.
- **Start Write Balancing** —See Starting and stopping thin pool write balancing on page 261.
- **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 (Enginuity 5773 - 5876)

This procedure explains how to create thin volumes on storage systems running Enginuity version 5773 - 5876. For instructions on creating thin volumes on storage systems running HYPERMAX OS 5977 or higher, refer to Creating thin volumes (HYPERMAX OS 5977) on page 217.

1. Select the storage system.
2. Hover the cursor over the **Storage** menu and select **Create Volumes** under **Common Tasks** to open the **Create Volume** dialog box.
3. In **Volume Type**, click **Virtual**.
4. Select **Configuration (TDEV or BCV + TDEV)** or thin volumes.
5. Select the **Emulation** type.
6. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
7. 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).

8. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **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:**
   - **None** — Allows the system to name the volumes (Default).
   - **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 231
2. To **Allocate Full Volume Capacity**, select the option.
3. If you selected to allocate capacity in the previous step, you can mark the allocation 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.
4. 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.
5. If Auto Meta is enabled on the system, and if you are attempting to create volumes larger than the **Minimum Meta Capacity**, specify values for the following in the **Define Meta** panel:
   - **Member capacity (Cyl/MB/GB)** — Size of the meta members to use when creating the meta volumes.
   - **Configuration (Striped/Concatenated)** — Whether to create striped or concatenated meta volumes.

6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Binding/Unbinding/Rebinding thin volumes

This procedure explains how to bind/unbind/rebind thin volumes to a thin pool of DATA 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 group level, requires Enginuity 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.
- For more information about thin pools and thin provisioning concepts, refer to the *Solutions Enabler Array Management CLI Product Guide*.

To bind/unbind/rebind thin volumes at the volume level:

1. Select the storage 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. *Optional*: Specify to **Allocate Full Volume capacity**.

   d. To view additional information on the selected volumes, click **Show selected volumes**.

   e. If you selected to allocate capacity in the previous step, you can mark the allocation 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.

   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. Click **OK** .

Rebinding:

a. Select one or more volumes and click >> 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 pool level:

1. Select the storage 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** wizard.

   b. Locate the volumes by selecting/typing values for any number of the following criteria:

      - **Capacity equal to** — Filters the list for volumes with a specific capacity.
      - **Volume ID** — Filters the list for a volume with specific ID.
      - **Volume Identifier Name** — Filters the list for the specified volume name.
      - **Additional Criteria** — Allows you to add conditional expressions to your query, as describe next:

        1. Select values for the following:

           - **Category** — Specifies the broad category.
           - **Attribute** — Refines the category.
           - **Operator** — Establishes the relationship between the attribute and the value.
           - **Value** — Conditions for the specified attribute.

        2. To add another expression, click **Add Another**, and repeat the previous step. To clear your criteria, click **Clear All**.

   c. Click **Find Volumes** to run the query. Results are displayed on the next page in the wizard.

   d. *Optional*: To modify your query, click **Modify Criteria** and repeat steps 5 and 6; otherwise continue with the remaining steps in this procedure.

   e. Select the volumes.

   f. *Optional*: Specify to **Allocate Full Volume capacity**, and then whether to mark the allocation 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.

   g. Click **Bind**.

Unbinding:

a. In **Related Objects**, click **Bound Volumes** to open the **Bound Volumes** 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. Click **OK**.

To bind/unbind/rebind thin volumes at the storage group level:

1. Select the storage 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. *Optional*: Specify to **Allocate Full Volume capacity**, and then whether to mark the allocation 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.

d. 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. Click **OK**.

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. Click **OK**.

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.

Note Virtual LUN migration requires Enginuity 5874 - 5876.

Virtual LUN Migration performs tiered storage migration by moving data from one RAID group to another, or from one thin pool to another. It is also fully interoperable with all other 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 Enginuity 5874 - 5876.
- 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 storage 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 long and are case sensitive.
5. Select a **Migration Type**. Available options are **Disk Group** and **Virtual Pool**.
6. Select a **Target Disk Group**.
7. Select the RAID **Protection type**.
8. Select **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** to create the migration session.

Migrating regular volumes

This procedure explains how to migrate individual regular volumes.

Before you begin:

- Virtual LUN migration requires Enginuity 5874 - 5876.
- To migrate volumes when access control is enabled, you must grant Symmetrix Optimizer access rights to the volumes.

To migrate regular volumes:

1. Select the storage 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 wether 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 - 5876.
- 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 storage 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. *Optional:* Expand **Show Advanced**. From the **Migrate allocations from pool** menu, select a pool from which to migrate allocations.
8. Click **OK**.

Migrating thin volumes

This procedure explains how to migrate individual thin volumes.

Before you begin:

- Virtual LUN migration requires Enginuity 5874 - 5876.
- 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 storage 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 wether to **Pin volumes** so that they cannot be moved by any FAST automated process.
8. *Optional:* Expand **Show Advanced**. From the **Migrate allocations from pool** menu, select a pool from which to migrate allocations.
9. Click **OK**.

Terminating a VLUN migration session

1. Select the storage system.
2. Select **Data Protection › Virtual LUN Migration** to open the **Virtual LUN 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 storage system.
2. Select **Data Protection › Virtual LUN Migration** to open the **Virtual LUN 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 the previous page
   - **View Details** — Viewing VLUN migration session details on page 487

External Storage

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.

At least four paths to external volumes is required, 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 storage 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** — Storage 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 storage ID.
- **Vendor** — External storage system 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 storage 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 293. (Only displays for Enginuity 5876)
- **Rescan** — See Rescanning external storage below.
- **Remove** —Removing external LUNs on page 288 (Only displays for HYPERMAX OS 5977 or higher)

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 storage 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.

Understanding ProtectPoint

ProtectPoint™ integrates your primary storage (VMAX system) with your protection storage (Data Domain system), thereby eliminating application servers from the data path. This integration reduces the cost and complexity associated with traditional backup applications while still providing the benefits of native backups.

Note TimeFinder/SnapVX operations are not supported on *working* ProtectPoint snapshots. TimeFinder/SnapVX operations are, however, supported to help repair failing ProtectPoint snapshots.

Protecting storage groups using ProtectPoint

This procedure explains how to backup thin volumes to Data Domain arrays.

Before you begin:

- The storage system must be running HYPERMAX OS 5977.
- You must have StorageAdmin permission.
- The Data Domain appliance must be connected and zoned to the storage system.
- Provide the Data Domain Admin the number and size of volumes in the production storage group and request that they provide you with double the number of similar volumes (masked/visible to the storage system). For example, if the production storage group contains 10 volumes, the Data Domain Admin should provide you with the LUN numbers of 20 similar volumes.

To protect storage groups using ProtectPoint:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.
5. Select **Backup Using ProtectPoint**.
6. Click **Next**.
7. Type the name of the **Point In Time Copy Name** and click **Next**.
8. Type a name for the **Backup Storage Group**, or leave the system-generated suggestion.
9. Select the external LUNs to add to the backup storage group and click **Add to Storage Group**. Note that the external LUNs you select must match in number and capacity the volumes in the production storage group.
10. Click **Next**.
11. Type a name for the **New Restore Storage Group**, or leave the system-generated suggestion.
12. Select the external LUNs to add to the restore storage group and click **Add to Storage Group**. Note that the external LUNs you select must match in number and capacity the volumes in the production storage group.
13. Click **Next**.
14. Verify your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
15. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
16. Once the job completes successfully, provide the following information to the Data Domain Admin:
    - The LUN numbers used in the backup storage group
    - The LUN numbers used in the restore storage group
    - The name of the point in time copy

Expanding ProtectPoint storage groups

This procedure explains how to increase the amount of storage in a storage group protected by ProtectPoint.

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- You must have StorageAdmin permission.
- The Data Domain appliance must be connected and zoned to the VMAX storage system.
- Provide the Data Domain Admin the number and size of volumes that you added to the production storage group and request that they provide you with double the number of similar volumes (masked/visible to the VMAX storage system). For example, if the production storage group contains 10 volumes, the Data Domain Admin should provide you with the LUN numbers of 20 similar volumes.

To expand protected storage groups:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Groups** panel, click **Total** to open the **Storage Groups** list view.
4. Select the storage group, click >>, and click **Expand ProtectPoint** to open the **Expand ProtectPoint** wizard.
5. Select the **Point In Time Copy**that you want to expand and click **Next**.
6. Select the external LUNs to add to the backup storage group and click **Add to Group**. Note that you need to select the same number of external LUNs as the number of volumes added to the production storage group.
7. Click **Next**.
8. Select the **Restore Storage Group**.
9. Select the external LUNs to add to the restore storage group and click **Add to Group**. Note that you need to select the same number of external LUNs as the number of volumes added to the production storage group.
10. Click **Next**.
11. Verify your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
12. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
13. Once the job has completed, provide the following information to the Data Domain Admin:
    - The LUN numbers added to the backup storage group
    - The LUN numbers added to the restore storage group
    - The name of the point in time copy

Removing external LUNs

This procedure explains how to remove external LUNs from storage groups protected with ProtectPoint.

Before you begin:

- This feature requires HYPERMAX OS 5977 or higher.
- LUNs must be virtualized.
- Encapsulated LUNs whose volumes are in a storage group cannot be removed.

To remove external LUNs:

1. Select the storage system.
2. Select **Storage › External Storage** to open the **External Storage** page.
3. *Optional*: 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.
4. From the filtered results table, select one or more LUNs and click **Remove** to open the **Remove External LUNs** dialog box.
5. *Optional*: To view details on the selected LUNs, click **Show selected external LUNs**.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Understanding FAST.X

FAST.X allows the seamless integration of VMAX3 Family storage systems and heterogeneous arrays. It enables LUNs on external storage to be used as raw capacity. Data services such as SRDF, TimeFinder, and Open Replicator are supported on the external device.

FAST.X requires HYPERMAX OS 5977, or higher.

For additional information on FAST.X, refer to the following documents:

- *Solutions Enabler Array Management CLI Guide*
- *Solutions Enabler TimeFinder CLI User Guide*
- *VMAX3 Family Product Guide*

Viewing external disks

The external disk list is available only for HYPERMAX OS 5977 or higher.

Note You must refresh the external disks list to view the latest status.

To view the external disks list using the Storage Groups Dashboard:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** section, select the SRP.
4. In the **Reports** section, click **View SRP Details**.
5. In the **Disk Groups** section, click the disk group name. The external disks list displays.

To view the external disks list using the Storage Resource Pools view:

1. Select the storage system.
2. Select **Storage › Storage Resource Pools** to open the **Storage Resource Pools** view.
3. Select the tab for the SRP.
4. In the **Disk Groups** section, click the disk group name. The external disks list displays.

Filtered results table

The following properties display:

- **External LUN WWN** — World Wide Name of the external disk.
- **Spindle** — Spindle ID of the external spindle.
- **Vendor** — Vendor name of the external disk.
- **Capacity (GB)** — Capacity in GB of the external disk.
- **Volume** — Volume ID on the external storage system.
- **Array ID**— ID of the storage system.
- **Service State** — Availability of the external disk. 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.
- **Disk State** — The state of the disk. Valid values are **Active**, **Drained**, **Draining**, and **Disabled**.
- **Drained** — Drain information about the disk if it is in **Drained** or **Draining** state. Otherwise it displays "-".

The following controls are available:

- **Add eDisks** — Adding external disks
- **Remove** — Removing external disks on the facing page
- **Start Draining** — Start draining external disks
- **Stop Draining** — Stop draining external disks
- **Activate** — Activating external disks

Adding external disks

You can add an external disk to the external disk group of a storage resource pool (SRP). You can complete the action from the SRP detailed view or from the external disk list view.

When adding an external disk for storage systems running HYPERMAX OS 5977 or higher, if there is no pre-existing external disk group, it is created automatically when the external disk is added to the selected SRP. If an external disk group exists for the external array's external LUN WWN, the external LUN WWN is added to it.

This action can be performed only for HYPERMAX OS 5977 or higher.

To add an external disk:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** section, select the SRP.
4. In the **Reports** section, click **View SRP Details**.
5. Do one of the following:
   - Click **Add eDisk**.

     Note If the SRP is an EMC CloudArray SRP and one or more external disks have already been added, further disks can be added only from the external disk list view.

   - In the **Disk Groups** panel, click the name of the disk group. In the external disk list view, click **Add eDisk**.

   The **Add eDisks** dialog box displays the available external LUN WWNs from multiple external arrays.
6. Select the external disk to be added.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Removing external disks

You can remove an external disk from a storage resource pool (SRP) if it is in a **Drained** state.

This action can be performed only for HYPERMAX OS 5977 or higher.

To remove an external disk:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** section, select the SRP from which you want to remove an external disk.
4. In the **Reports** section, click **View SRP Details**.
5. In the **Disk Groups** section, click the disk group name. The external disks list displays.
6. Select the external disk that you want to remove and click **Remove**. The **Remove External LUNs** dialog box displays and prompts for confirmation that you want to remove the external disk.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Start draining external disks

Before you begin:

The storage resource pool (SRP) containing the external disk you want to drain must have sufficient free space to absorb the allocated tracks from the external disk that is being drained.

You can drain a disk only if it is not currently draining or already drained.

To start a drain operation on an external disk:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** section, select the SRP.
4. In the **Reports** section, click **View SRP Details**.
5. In the **Disk Groups** section, click the disk group name. The external disks list displays.
6. Click **Start Draining**. The **Drain eDisk Dialog** displays and prompts for confirmation that you want to start a draining operation.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Stop draining external disks

Before you begin:

You can stop the drain operation on an external disk only if it is currently draining.

To stop a draining operation on an external disk:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** section, select the SRP.
4. In the **Reports** section, click **View SRP Details**.
5. In the **Disk Groups** section, click the disk group name. The external disks list displays.
6. Click **Stop Draining**. The **Drain eDisk Dialog** displays and prompts for confirmation that you want to stop a draining operation.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Activating external disks

You can activate an external disk if it is in a draining, drained, or disabled state.

Before you begin:

This action can be performed only for HYPERMAX OS 5977 or higher.

To activate an external disk:

1. Select the storage system.
2. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
3. In the **Storage Resource Pools** section, select the SRP.
4. In the **Reports** section, click **View SRP Details**.
5. In the **Disk Groups** section, click the disk group name. The external disks list displays.
6. Select one or more external disks. Click **Activate**.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Understanding Federated Tiered Storage

Federated Tiered Storage (FTS) allows you to attach external storage to a VMAX system. Attaching external storage allows you to use physical disk space on existing VMAX systems while gaining access to VMAX features such as local replication, remote replication, storage tiering, data management, and data migration.

FTS requires Enginuity 5876 and is not supported on VMAX systems running HYPERMAX OS 5977 or higher.

For additional information on FTS, refer to the following documents:

- *Symmetrix Federated Tiered Storage (FTS) Technical Notes*
- *Solutions Enabler Array Management CLI Product Guide*
- *Solutions Enabler TimeFinder Family CLI User Guide*
- *Symmetrix VMAX Family with Enginuity (10K, 20K, 40K) Product Guide*

Virtualizing external LUNs

When you attach external storage to a Symmetrix system, FAST.X 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.

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

Geometry of encapsulated volumes

Enginuity builds storage volumes based on the storage system cylinder size (fifteen 64 K tracks), so the capacity of storage 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 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 *Solutions Enabler Array Controls CLI Guide*.

To virtualize external LUNs:

1. Select the storage 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.

   WARNING 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 851 and Previewing on page 850.
   - **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 storage 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 851 and Previewing on page 850.

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 - 5876 and storage groups.

Creating storage templates

This procedure explains how to create storage templates.

Before you begin:

- Storage templates require Enginuity 5874 - 5876.
- 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 storage 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. Select **Create BCV Volume**, if required.

   c. If the template will be associated with a specific storage 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 preallocate the entire volume, select **Allocate Full Volume Capacity**.

   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 **Create Meta** and complete the following steps:

   a. Select the **Meta Volume Configuration**. The storage template uses only meta volumes of the specified configuration.

   b. Select the **Meta Volume 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**.

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 - 5876.
- 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\EMC\SMAS\jboss\domain\servers\server-0\data\Template Files**).

Importing storage templates

This procedure explains how to import storage templates to a Unisphere server.

Before you begin:

- Storage templates require Enginuity 5874 - 5876.
- 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 - 5876.
- 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 **Delete** in the confirmation box.

Viewing storage templates

1. Select the storage 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** — Name of the template.
   - **Symmetrix ID** — Storage 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 (GB)** — 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 296.
   - **View Details** — See Viewing storage template details on the next page.
   - **Import** — See Importing storage templates on the previous page.
   - **Export** — See Exporting storage templates on the previous page.
   - **Delete** — See Deleting storage templates on the previous page.

Viewing storage template details

1. Select the storage 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.

   Depending on configuration, some of 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 storage system. Any indicates that the template is associated with all supported storage 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 Volume Member Capacity** — Capacity of each meta member.
   - **Meta Volume 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.
   - **Disk Technology** — Disk technology type.
   - **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 296.
- **Delete** — See Deleting storage templates on page 298.
- **Apply** — Applies changes made in the **Properties** list. For example, renaming the storage group.
- **Cancel** — Cancels changes made in the **Properties** list.

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 reserved 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 storage system. For more information, refer to See Setting system preferences on page 38.

To reserve volumes:

1. Select the storage system.
2. Select **System › Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **View Reservations** to open the **Reservations** list view.
4. Click **Create** to open the **Create Reservation** wizard.
5. Locate the volumes by selecting/typing values for any number of the following criteria:
   - **Capacity equal to** — Filters the list for volumes with a specific capacity.
   - **Volume ID** — Filters the list for a volume with specific ID.
   - **Volume Identifier Name** — Filters the list for the specified volume name.
   - **Additional Criteria** — Allows you to add conditional expressions to your query, as describe next:
     a. Select values for the following:
        - **Category** — Specifies the broad category.
        - **Attribute** — Refines the category.
        - **Is equal to** — Operator that establishes the relationship between the attribute and the value.
        - **Value** — Conditions for the specified attribute.
     b. To add another expression, click **Add Another**, and repeat the previous step. To clear your criteria, click **Clear All**.
6. Click **Find Volumes**.
7. Select the **Volumes** .
8. Type any **Comments** regarding the reservation. The requirement for this field is set in system preferences.

9. Select **Expiration** and click calendar to open the **Reservation Expiration** dialog box.
10. Select an **Expiration Date** and **Expiration Time**.
11. Click **OK** to close the **Reservation Expiration** dialog box.
12. Click **Reserve** to create the reservation.

Adding volumes to reservations

1. Select the storage system.
2. Select **System › Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **View Reservations** to open the **Reservations** list view.
4. Select the reservation and click **View Details** to open its **Details** view.
5. Click **Add** to open the **Add Volumes to Reservations** wizard.
6. Locate the volumes by selecting/typing values for any number of the following criteria:
   - **Capacity equal to** — Filters the list for volumes with a specific capacity.
   - **Volume ID** — Filters the list for a volume with specific ID.
   - **Volume Identifier Name** — Filters the list for the specified volume name.
7. Click **Find Volumes** to run the query now, or **Show Advanced** to add conditional expressions to your query, as described next.
   a. Select values for the following, and then click **Add**.
      - **Category** — Specifies the broad category.
      - **Attribute** — Refines the category.
      - **Operator** — Establishes the relationship between the attribute and the value.
      - **Value** — Conditions for the specified attribute.

      b. Repeat the previous step for each additional expression.

   c. Click **Find Volumes**.
8. Select the **Volumes** and click **Add to Reservation**.

Removing volumes from reservations

1. Select the storage system.
2. Select **System › Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **View Reservations** to open the **Reservations** list view.
4. Select a reservation and click **View Details**.
5. In the **Related Objects** section of the **Properties** dialog, click **Reserved Volumes**.
6. In the Reserved Volumes list, select the reservation and click **Remove**.
7. Click **OK** in the confirmation message.

Releasing reservations

1. Select the storage system.
2. Select **System › Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **View Reservations** to open the **Reservations** list view.
4. Select one or more reservations and click **Release**.
5. Click **OK** in the confirmation message.

Viewing reservations

1. Select the storage system.
2. Select **System › Dashboard** to open the **System Dashboard**.
3. In the **Summary** panel, click **View 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 page 302
- **View Detail** — See Viewing reservation details below.
- **Release** — Releasing reservations above

Viewing reservation details

1. Select the storage system.
2. Select **System › Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **View Reservations** to open the **Reservations** list view.
4. 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 302
- **Release** — Releasing reservations on the previous page
- **Add** — Adding volumes to reservations on page 303

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.

Understanding Optimizer

Optimizer improves storage system 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.

Note Optimizer is supported on storage systems running Enginuity version 5671 - 5876.

For information on Optimizer concepts, refer to the *VMAX Product Guide*.

Managing Optimizer

The **Optimizer** dashboard provides you with a single place from which to manage Optimizer.

Before you begin:

Optimizer is supported on storage systems running an Enginuity version lower than 5977.

To access the Optimizer dashboard:

1. Select the storage 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 one week.
  - **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, with the default being eight hours.

- **Optimizer Specific Settings:**
  - **Startup Mode** — Specifies whether Optimizer starts automatically with the 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 310.
- **Swap/Move History** — Viewing Optimizer swap/move history on page 310
- **Approve Swap** — See Approving Optimizer Swaps on page 309.
- **Start/Lock** — Lock/Unlock/Start/Stop dialog box (Does not apply/display for Enginuity 5874 or higher).
- **Enable/Disable** — Enabling/Disabling Optimizer on the 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:

- Optimizer is supported on storage systems running an Enginuity version lower than 5977.

- This procedure is only required when performing operations on a storage system running an Enginuity version lower than 5874.

To start/stop Optimizer:

1. Select the storage 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

Before you begin:

- Optimizer is supported on storage systems running an Enginuity version lower than 5977.
- This procedure explains how to enable/disable Optimizer on a Symmetrix system running Enginuity 5874 or higher:

To enable/disable Optimizer:

1. Select the storage 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 **Close** in the confirmation message.

Locking/Unlocking Optimizer

The following explains how to lock/unlock Optimizer controls. Locking Optimizer prevents other uses from making changes while you are working with the software.

Before you begin:

- Optimizer is supported on storage systems running an Enginuity version lower than 5977.
- This procedure is only required when performing operations on a storage system running an Enginuity version lower than 5874.

To start/stop Optimizer:

1. Select the storage 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

Before you begin:

Optimizer is supported on storage systems running an Enginuity version lower than 5977.

To approve Optimizer swaps:

1. Select the storage 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 storage 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.

   Enginuity 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**. Do one of the following:
   - Select **According to Optimizer policy**.

- Specify a date and time.
- Select **Decline Plan**.

Viewing Optimizer swap/move lists

The following explains how to view information on the current Optimizer/FAST plan:

1. Select the storage 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 Enginuity version running on the storage system:

   Enginity 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.

   Enginuity 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.

4. *Optional:* To refresh the list of swaps/moves, click **Refresh**.

Viewing Optimizer swap/move history

The following explains how to view information on the past Optimizer/FAST swaps/moves:

1. Select the storage 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 Enginuity version running on the storage system:

   Enginity 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.

   Enginuity 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).

Managing vVols

The **VVols Dashboard** provides you with a single place to monitor and manage vVols.

Before you begin:

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

To access the VVols Dashboard:

1. Select the storage system.
2. Select **Storage › VVols Dashboard**.

VVols Dashboard:

The **VVols Dashboard** is organized into the following panels:

Summary

Displays the following vVol summary information:

- **Storage Containers**
  The number of storage containers on the selected storage system.
- **Protocol Endpoints**
  The number of protocol endpoints on the selected storage system.
- **PE Masking Views**
  The number of masking views that contain protocol endpoints.

To view additional information on a particular object, click its icon to open the corresponding list view.

Common Tasks

Displays links to the following common tasks:

- **Create Storage Container** — Creating storage containers on page 314
- **Provision Protocol Endpoint to Host** — Provisioning protocol endpoints to hosts on page 318
- **Storage Container Alerts** — Viewing storage container alerts

VASA Provider Status

Displays an icon (or) representing the status of the VASA provider.

To refresh the status of the VASA provider, click .

Symmetrix Consumed Capacity - Subscribed

Displays a bar graph representing how much subscribed space all storage containers consume on the storage system.

Storage Resources Consumed Capacity - Subscribed

Displays a list of storage resources within all containers on the storage system, showing the current usage of each storage resource, ascending by usage.

- **Name** — The name of the capability profile.
- **Subscribed Used(%)** — The current percent of subscribed tracks within the storage resource in relation to the limit imposed on the capability profile.
- **Limit (GB)** — The subscribed limit imposed on the storage resource.
- **Container** — The name of the storage container with which the storage resource is associated.

Viewing storage containers

To view the storage container list:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.

   The following properties display:

   - **Name** — The name of the storage container.
   - **Storage Resources** — The number of associated storage resources.
   - **Subscribed Used (%)** — The current percentage of subscribed tracks within the storage container, in relation to the limit imposed on all of the storage resources within the storage container.
   - **Subscribed Limit (GB)** —The current total limit of all storage resources in GB.

   The following controls are available:

   - **Create Container** — Creating storage containers on the next page
   - **Delete** — Deleting storage containers on page 315
   - **View Details** — Viewing storage resource details below

Viewing storage resource details

To view storage resource details:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container. Click **View Details**.

   The following properties display:

   - **Name** — The name of the storage container.
   - **Description** — The description of the storage container. This field is editable.

- **Subscribed Capacity Limit (GB)** — The total combined limit of all storage resources within the storage container.
- **Subscribed Capacity Used (GB)** — The current subscribed usage on the storage container of all of the storage resources within the storage container.
- **Subscribed Capacity Free (GB)** — The total free subscribed capacity, based on the capacity used and the limit of all of the storage resources in the storage container.
- **Number of Storage Resource** — The total number of storage resources within the storage container.

The following controls are available:

- **Delete** — Deleting storage containers on the facing page
- **Apply** — Applies any changes made to the storage container details.
- **Cancel** — Discards any changes made to the storage container details.

Creating storage containers

To create a storage container:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Click **Create Container**. The **Create Storage Container** wizard displays.
5. On the **Storage Container** page, complete the following steps:
   a. Type a name for the storage container.
   b. *(Optional)* Type a description of a storage container.
6. Click **Next**.
7. On the **Storage Resources** page, specify at least one storage resource. Default values for a new storage resource are populated.

   To remove a storage resource from the list of associated storage resources, hover the mouse over the storage resource and click ✕.

   To add a storage resource, click **Add Resource** and complete the following steps:
   a. In the **Name** field, type a name for the storage resource, or accept the default name.
   b. From the **SRP Name** menu, select the SRP to apply to the storage resource.
   c. From the **SLO Name** menu, select the SLO to apply to the storage resource.
   d. From the **Workload** menu, select the workload to apply to the storage resource.
      To view a popup showing the current workload demand report click ℹ.
   e. *(Optional)* In the **Subscribed Limit (GB)** field, type the imposed subscribed limit on the storage resource.
8. Click **Next**.
9. On the **Review** page, review the details and do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Deleting storage containers

To delete a storage container:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container you want to delete and click **Delete**.
5. Click **OK** in the confirmation message.

Viewing storage resources

To view the storage resource list:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container. Click **View Details**.
5. In the **Related Objects** pane, click **Storage Resources** to display the **Storage Resources** list view.

   The following properties display:

   - **Name** — The name of the capability profile.
   - **SRP** — The number of the SRP.
   - **SLO** — The name of the SLO.
   - **Workload** — The name of the workload.
   - **Used (%)**— The current percent of subscribed tracks within the storage resource in relation to the limit imposed on the storage resource.
   - **Subscribed Used (GB)** — The capacity used within the storage resource.

   The following controls are available:

   - **Add** — Adding storage resources to storage containers on page 317
   - **Remove** — Removing storage resources from storage containers on page 317
   - **View Details** — Discards any changes made to the storage container details

Viewing storage resource details

To view a storage resource details:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container. Click **View Details**.
5. In the **Related Objects** pane, click **Storage Resources** to display the **Storage Resources** list view.
6. Select the storage resource. Click **View Details**.

   The following properties display:

   - **Name** — The name of the storage resource.
   - **Storage Container** — The name of the associated storage container.
   - **SRP** — The name of the associated SRP.
   - **SLO** — The name of the associated SLO.
   - **Workload** — The name of the associated workload.
   - **Subscribed Capacity Limit (GB)** — The subscribed capacity limit imposed. This field is editable.
   - **Subscribed Capacity Used (GB)** — The current subscribed usage on the storage resource.
   - **Subscribed Capacity Free (GB)** — The subscribed free space on the storage resource.

   The following controls are available:

- **Remove** — Removing storage resources from storage containers on the facing page
- **Apply** — Applies any changes made to the storage container details
- **Cancel** — Discards any changes made to the storage container details

Viewing storage resource related SRPs

To view the related SRPs of a storage resource:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container. Click **View Details**.
5. In the **Related Objects** pane, click **Storage Resources** to display the **Storage Resources** list view.
6. Select the storage resource. Click **View Details**.

7. In the **Related Objects** pane, click the SRP link to display the **Storage Resource Pools** details view.
   For more information about the **Storage Resource Pools** list view, refer to Viewing Storage Resource Pools.

Adding storage resources to storage containers

To add a storage resource to a storage container:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container. Click **View Details**.
5. In the **Related Objects** pane, click **Storage Resources**.
6. Click **Add**.
7. On the **Review** page, review the details and do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Removing storage resources from storage containers

To remove a storage resource from a storage container:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container. Click **View Details**.
5. In the **Related Objects** pane, click **Storage Resources**.
6. Select the storage resource you want to remove and click **Remove**.
7. In the **Remove Storage Resource** dialog, click **OK**.

Viewing protocol endpoints

To view the protocol endpoints list:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Protocol Endpoints** to display the **Protocol Endpoints** list view.

   The following properties display:

   - **Name** — The volume ID of the protocol endpoint.

- **Masking View**— Indicates, using a ✅ or ❌ symbol, if the protocol endpoint is in a masking view or not.
- **Storage Groups** — The number of associated storage groups.
- **Reserved** — Indicates if the protocol endpoint is reserved or not.

The following controls are available:

- **Provision PE to Host** — Provisioning protocol endpoints to hosts below
- **Set Volume Identifiers** — Setting volume identifiers on page 231.
- **View Details** — Viewing storage resource details on page 313

Viewing protocol endpoint details

To view protocol endpoint details:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. Click **Protocol Endpoints** to display the **Protocol Endpoints** list view.
4. Select the protocol endpoint. Click **View Details**.

   The following properties display:

   - **Name** — The name of the protocol endpoint.
   - **Volume Identifier** — The volume identifier of the protocol endpoint.
   - **Status** — The status of the protocol endpoint.
   - **Reserved** — The reserved status of the protocol endpoint. Valid values are **Yes** and **No**.
   - **Number of Storage Groups** — The total number of storage groups associated with the protocol endpoint.
   - **Number of Masking Views** — The total number of masking views associated with the protocol endpoint.
   - **Number of Bound vVols**— The total number of vVols bound with the protocol endpoint.

Provisioning protocol endpoints to hosts

To provision a protocol endpoint to a host:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. In the Common Tasks pane, click **Provision Protocol Endpoint to Host** to display the **Provision Protocol Endpoint to Host** wizard.
4. On the **Select Host/Host Group** pane, specify a host or host group. Do one of the following:
   - Select an existing host or host group from the list.

- To create a new host, click **Create Host**. The **Create Host** dialog displays. For more information, refer to Creating hosts.
- To create a host group, click **Create Host Group**. The **Create Host Group** dialog displays. For more information, refer to Creating host groups.

5. Click **Next**.
6. On the **Select Port Group** pane, specify a port group. Do one of the following:
   - To create a new port group, select **New**.
     For more information about creating port groups, refer to Creating port groups on page 350
   - To use an existing port group, select **Existing**, and select a port group from the **Port Group** list.
7. Click **Next**.
8. On the **Review** page, review the details and do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing storage container alerts

To view the storage container alert list:

1. Select the storage system.
2. Select **Storage › VVols Dashboard** to open the **VVols Dashboard**.
3. In Common Tasks pane, click **Storage Container Alerts**.
4. *Optional*: Use the alert filter to view a subset of the listed alerts. For more information on the Alert Filter, refer to Filtering alerts (page 46).

   The following properties display:

   - **State** — State of the alert. Possible values are New or Acknowledged.
   - **Severity** — Alert severity. Possible values are:
     - Fatal
     - Critical
     - Warning — The following events map to this severity:
       - The component is in a degraded state of operation.
       - The storage array is no longer present (during certain operations).
       - The component is in an unknown state.
       - The component is (where possible) in a write-disabled state.
     - Information — The component is no longer present (during certain operations).
     - Normal — The component is now (back) in a normal state of operation.

- **Type** — Type of alert. Possible values are Array, Performance, Server, System, and File.
- **Object** — Component to which the alert is related. This is because server alerts are specific to the server or runtime environment and are not associated with a specific object or storage system.
- **Description** — Description of the alert.
- **Created** — Date/time the alert was created.
- **Acknowledged** — Date/time the alert was acknowledged.

The following controls are available:

- **View Details** — See Viewing alert details on page 46.
- **Acknowledge** — See Acknowledging alerts on page 46.
- **Delete** — See Deleting alerts on page 47.

CHAPTER 5
Host Management

This chapter contains the following:

Hosts 322
Host Groups 331
Masking views 338
Initiators 343
Port groups 350
Virtual servers 357
CU images 362
Host aliases 369
Host Cache Adapters 373

Hosts

Creating hosts

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 - 5876, or HYPERMAX OS 5977 or higher.
- The maximum number of initiators allowed in a host depends on the storage operating environment:
  - For Enginuity 5874 - 5876, the maximum allowed is 32.
  - For HYPERMAX OS 5977 or higher, the maximum allowed is 64.

To create hosts:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Click **Create Host** to open the **Create Host** dialog.
4. Type a **Host Name**. Host names must be unique from other hosts/host groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host names are case-insensitive.
5. Select the **Fibre** radio button to filter the available initiators table to display Fibre Channel initiators only or select the **ISCSI** radio button to filter the table to display iSCSI initiators only. The **Fibre** radio button is selected by default.
6. Add initiators to the host:
   a. Specify the initiator by typing its name or by selecting it from the list. To filter the list, type part of a initiator name. Repeat this step for each additional initiator.

   Note Initiators can only belong to one host at a time; therefore, any initiators that do not appear in the list already belong to another host.

b. Click **Add**.

c. Repeat these steps for each additional initiator.

7. *Optional*: Set the host port attributes:

   a. Click **Set Host Flags** to open the **Set Host/Host Group Flags** dialog box.

   b. *Optional*: Select a host whose flag settings you want to copy.

   c. 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 host must be kept consistent for all volumes within each masking view of which this host is part. When set, any masking operation involving this host that would result in inconsistent LUN values, will be rejected. When not set, the storage 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 Host/Host Group Flags** dialog box.

8. 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 851.

Adding initiators to hosts

This procedure explains how to add host initiators.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.
- On storage systems running HYPERMAX OS 5977 or higher, iSCSI and fibre initiators cannot be mixed in the same host.

To add initiators to hosts:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Initiators** to open the host's **Initiators** list view.
5. Click **Add Initiators** to open the **Add Initiators** dialog box.
6. Specify the initiator by typing its name or by selecting it from the list. The Initiators table is a filtered list based on whether the initiator is Fibre Channel or iSCSI.To filter the list, type part of the initiator name. Click **Add**. Repeat this step for each additional host.
7. Click **OK**.

Removing initiators from hosts

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

To remove initiators from hosts:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Initiators** to open the host's **Initiators** list view.
5. Select the initiator and click **Remove**.
6. Click **OK** in the confirmation message.

Modifying hosts

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

Modifying hosts:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host, click more >> , and select **Modify** to open the **Modify Host** dialog box.
4. To change the **Host Name**, highlight it and type a new name over it. Host names must be unique from other hosts on the Symmetrix system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host names are case-insensitive.

5. Do the following, depending on whether you are adding or removing initiators:
6. **Adding initiators:**
   a. In the **Select Initiators** list box, type the initiator name or select it from the list. To filter the list, type part of a initiator name.

      Note Initiators can only belong to one host at a time; therefore, any initiators that do not appear in the list already belong to another host.
      The **Add Initiators** table is a filtered list based on whether the host is Fibre Channel or iSCSI.

   b. Click **Add**.
   c. Repeat these steps for each additional initiator.

   Removing initiators:

   a. In the **Initiators Added to Host** list, select the initiator and click **Remove**.
   b. Click **OK** in the confirmation message.
   c. Repeat these steps for each additional initiator.
7. 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 851.

Renaming hosts/host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

To rename host/host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group from the list and click **View Details** to open its **Details** view.
4. In the **Properties** panel, type a new name for the host/host group and click **Apply**. Host/host group names must be unique from other hosts/host groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host/host group names are case-insensitive.

Setting host/host group flags

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group, click more >> , and select **Set Flags** to open the **Set Host/Host Group Flags** dialog box.
4. *Optional*: Select a host/host group whose flag settings you want to copy.
5. Modify any of the flags, 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 host must be kept consistent for all volumes within each masking view of which this host is part. When set, any masking operation involving this host that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations.
7. Click **OK**.

Deleting hosts/host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

Deleting hosts/host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group, click more >> , and select **Delete**.
4. Click **Delete** in the confirmation dialog box.

Viewing hosts/host groups

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.

   Use the **Hosts** list view to view and manage hosts.

   The following properties display:

- **Name** — Host/host group name. An arrow icon at the beginning of the name indicates that the host is a host group. Click the icon to view hosts contained in the group.
- **Masking Views** — Number of masking view associated with the host.
- **Initiators** — Number of initiators in the host.
- **Consistent LUNs** — Flag indicating if the Consistent LUNs flag is set. When set, any masking operation involving this host/host group that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations. ✓ indicates that the feature is set.
- **Port Flag Overrides** — Flag indicating if any port flags are overridden for the host. ✓ indicates that there are overridden port flags.
- **Last Update** — Timestamp of the most recent changes to the host.

The following controls are available:

- **Create Host** — See Creating hosts on page 322.
- **Create Host Group** — Creating host groups on page 331
- **Provision Storage to Host** — Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106/Using the Provision Storage wizard (Enginuity 5874 - 5876) on page 110
- **View Details** — See Viewing host/host group details on page 336.
- **Modify** — Modifying hosts on page 324 /Modifying host groups on page 333
- **Set Flags** —Setting host/host group flags on page 334
- **Delete** — Deleting hosts/host groups on page 334

Viewing host/host group details

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group and click **View Details** to open its **Details** view.

   Use the host **Details** view to view and manage hosts/host groups. This view contains two panels, the Properties and Related Objects.

Properties panel

Note The properties and controls available in this panel depend on whether you are viewing details of an individual host or of host group, and on the storage operating environment.

The following properties display:

- **Name** — Host/host group name. To rename the host/host group, type a new name over the existing and click **Apply**. Host/host group names must be unique from other hosts/host groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host names are case-insensitive.
- **Number of Hosts** — Number of hosts in the group. This field only displays for host groups.
- **Number of Host Groups** — Number of host groups in which this host is a member. This field only displays for individual hosts.
- **Number of Initiators** — Number of initiators in the host/host group. For host groups, the value includes initiators in any child host groups.
- **Number of Masking Views** — Number of masking views with which the host/host group is associated.
- **Consistent LUNs** — Flag indicating if the Consistent LUNs flag is set. When set, any masking operation involving this host/host group that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations. ✓ indicates that the feature is set.
- **Port Flag Overrides** — Flag indicating if any port flags are overridden for the host. ✓ indicates that there are overridden port flags.
- **Enabled Port Flags** — List of any enabled port flags overridden by the host/host group.
- **Disabled Port Flags** — List of any disabled port flags overridden by the host/host group.
- **Last Update** — Timestamp of the most recent changes to the host/host group.

The following controls are available:

- **Provision Storage to Host** —Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106/Using the Provision Storage wizard (Enginuity 5874 - 5876) on page 110
- **Modify** —Modifying hosts on page 324/Modifying host groups on page 333
- **Set flags** — See Setting host/host group flags on page 334.
- **Delete** — See Deleting hosts/host groups on page 334.
- **Apply** — Applies new host 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 host initiators

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host and click **View Details** to open its **Details** view.
4. From the **Related Objects** panel, click **Initiators** to open the host **Initiators** list view.

   The following properties display:

   - **Initiator** — 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 Initiator** — Adding initiators to hosts on page 323
   - **Remove** — Removing initiators from hosts on page 324
   - **View Details** — See Viewing initiator details on page 347.

Host/Host group flags

Table 4: Host/Host group flags

Attribute	Description
Common Serial Number	Enables a unique serial number. This attribute is only available on storage systems running Enginuity 5773 - 5876.
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 (storage system 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.
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.
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).

Attribute	Description
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.
AS400	Indicates whether AS/400 is enabled. This attribute is only available on storage system's running Enginuity 5773 - 5876.
Open VMS*,**	Enables an Open VMS fiber connection.

* To enable/disable this flag when it is already overridden (i.e., the **Override** option is already selected), you must:

1. Clear the **Override** option and click **OK** to close the dialog.
2. Open the dialog again, select **Override**, and then the desired state (**Enable**/**Disable**).
3. Click **OK**.

** For storage systems running HYPERMAX OS 5977 or higher, if Volume Set Addressing is overridden and enabled, the Open VMS flag must be disabled. However, if you do not actually select the Open VMS override option, Solutions Enabler will override and disable it.

If the Open VMS flag is overridden and enabled, the Volume Set Addressing flag must be disabled. However, if you do not actually select the Volume Set Addressing override option, Solutions Enabler will automatically override and disable it.

Host Groups

Creating host groups

This procedure explains how to create a host group (collection of hosts). For instructions on creating a host, refer to Creating hosts on page 322.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 - 5876, or HYPERMAX OS 5977 or higher.
- The maximum number of hosts allowed in a host group depends on the HYPERMAX OS:
  - For Enginuity 5874 - 5876, the maximum allowed is 32.
  - For HYPERMAX OS 5977 or higher, the maximum allowed is 64.

To create host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Click **Create Host Group** to open the **Create Host Group** dialog box.
4. Type a **Host Group Name**. Host group names must be unique from other hosts/host groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host group names are case-insensitive.
5. Select the **Fibre** radio button to filter the available hosts table to display Fibre Channel hosts only or select the **ISCSI** radio button to filter the table to display iSCSI hosts only. The **Fibre** radio button is slected by default.
6. Add hosts to the group:
   - To create new hosts to add to the group, click **Create New Host** to open the **Create New Host** dialog box. For instructions on creating hosts, refer to Creating hosts on page 322.
   - To add existing hosts to the group:
     a. Specify the host by typing its name or by selecting it from the list. To filter the list, type part of the host name. Repeat this step for each additional host.
     b. Click **Add**.
     c. Repeat these steps for each additional host.
7. *Optional*: Set the host port attributes:
   a. Click **Set Host Flags** to open the **Set Host/Host Group Flags** dialog box.
   b. *Optional*: Select a host whose flag settings you want to copy.
   c. 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 host must be kept consistent for all volumes within each masking view of which this host is part. When set, any masking operation involving this host that would result in inconsistent LUN values, will be rejected. When not set, the storage 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 Host/Host Group Flags** dialog box.

8. 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 851.

Adding hosts to host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 - 5876, or HYPERVMAX OS 5977 or higher.

Adding hosts to host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host group (or empty host) and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Hosts** to open the group's **Hosts** list view.
5. Click **Add Hosts** to open the **Add Host** dialog box.
6. Specify the host by typing its name or by selecting it from the list. The hosts table is a filtered list based on whether the host selected is Fibre Channel or ISCSI.To filter the list, type part of the host name. Click **Add**. Repeat this step for each additional host.
7. Click **OK**.

Removing hosts from host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The VMAX array must be running Enginuity version 5874 or higher.

Modifying host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host group and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Hosts** to open the group's **Hosts** list view.
5. Select the host and click **Remove**.
6. Click **OK** in the confirmation message.

Modifying host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

Modifying host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host group, click more >> , and select **Modify** to open the **Modify Host Group** dialog box.
4. To change the host group **Name**, highlight it and type a new name over it. Host names must be unique from other hosts on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host names are case-insensitive.
5. Do the following, depending on whether you are adding or removing hosts:
6. **Adding Host:**
   a. *Optional*: To create a new host to add to the group, click **Create Host** to open the **Create Host** dialog box. For help, refer to Creating hosts on page 322.
   b. Select the host and click **Add**. To filter the list, type part of the host name. Repeat this step for each additional host.

   Removing Hosts:
   a. In the list of hosts, select the host and click **Remove**.
   b. Click **OK** in the confirmation message.
   c. Repeat these steps for each additional host.
7. 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 851.

Renaming hosts/host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

To rename host/host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group from the list and click **View Details** to open its **Details** view.
4. In the **Properties** panel, type a new name for the host/host group and click **Apply**. Host/host group names must be unique from other hosts/host groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host/host group names are case-insensitive.

Setting host/host group flags

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group, click more >> , and select **Set Flags** to open the **Set Host/Host Group Flags** dialog box.
4. *Optional*: Select a host/host group whose flag settings you want to copy.
5. Modify any of the flags, 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 host must be kept consistent for all volumes within each masking view of which this host is part. When set, any masking operation involving this host that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations.
7. Click **OK**.

Deleting hosts/host groups

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5874 or higher.

Deleting hosts/host groups:

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group, click more >> , and select **Delete**.
4. Click **Delete** in the confirmation dialog box.

Viewing hosts/host groups

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.

   Use the **Hosts** list view to view and manage hosts.

   The following properties display:

   - **Name** — Host/host group name. An arrow icon at the beginning of the name indicates that the host is a host group. Click the icon to view hosts contained in the group.
   - **Masking Views** — Number of masking view associated with the host.
   - **Initiators** — Number of initiators in the host.
   - **Consistent LUNs** — Flag indicating if the Consistent LUNs flag is set. When set, any masking operation involving this host/host group that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations. ✓ indicates that the feature is set.
   - **Port Flag Overrides** — Flag indicating if any port flags are overridden for the host. ✓ indicates that there are overridden port flags.
   - **Last Update** — Timestamp of the most recent changes to the host.

   The following controls are available:

   - **Create Host** — See Creating hosts on page 322.
   - **Create Host Group** — Creating host groups on page 331
   - **Provision Storage to Host** — Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106/Using the Provision Storage wizard (Enginuity 5874 - 5876) on page 110
   - **View Details** — See Viewing host/host group details on the next page.
   - **Modify** — Modifying hosts on page 324 /Modifying host groups on page 333
   - **Set Flags** —Setting host/host group flags on the previous page
   - **Delete** — Deleting hosts/host groups on the previous page

Viewing host/host group details

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host/host group and click **View Details** to open its **Details** view.

   Use the host **Details** view to view and manage hosts/host groups. This view contains two panels, the Properties and Related Objects.

Properties panel

Note The properties and controls available in this panel depend on whether you are viewing details of an individual host or of host group, and on the storage operating environment.

The following properties display:

- **Name** — Host/host group name. To rename the host/host group, type a new name over the existing and click **Apply**. Host/host group names must be unique from other hosts/host groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Host names are case-insensitive.
- **Number of Hosts** — Number of hosts in the group. This field only displays for host groups.
- **Number of Host Groups** — Number of host groups in which this host is a member. This field only displays for individual hosts.
- **Number of Initiators** — Number of initiators in the host/host group. For host groups, the value includes initiators in any child host groups.
- **Number of Masking Views** — Number of masking views with which the host/host group is associated.
- **Consistent LUNs** — Flag indicating if the Consistent LUNs flag is set. When set, any masking operation involving this host/host group that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations. ✓ indicates that the feature is set.
- **Port Flag Overrides** — Flag indicating if any port flags are overridden for the host. ✓ indicates that there are overridden port flags.
- **Enabled Port Flags** — List of any enabled port flags overridden by the host/host group.
- **Disabled Port Flags** — List of any disabled port flags overridden by the host/host group.
- **Last Update** — Timestamp of the most recent changes to the host/host group.

The following controls are available:

- **Provision Storage to Host** —Using the Provision Storage wizard (HYPERMAX OS 5977 or higher) on page 106/Using the Provision Storage wizard (Enginuity 5874 - 5876) on page 110
- **Modify** —Modifying hosts on page 324/Modifying host groups on page 333
- **Set flags** — See Setting host/host group flags on page 334.

- **Delete** — See Deleting hosts/host groups on page 334.
- **Apply** — Applies new host 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 hosts in host groups

1. Select the storage system.
2. Select **Hosts › Hosts** to open the **Hosts** list view.
3. Select the host group and click **View Details** to open its **Details** view.
4. In the **Related Objects** panel, click **Hosts** to open the host group - **Hosts** list view.

   Use the host group - **Hosts** list view to view and manage hosts in the host group.

   The following properties display:

   - **Name** — Host/host group name.
   - **Masking Views** — Number of masking view associated with the host.
   - **Initiators** — Number of initiators in the host.
   - **Consistent LUNs** — Flag indicating if the Consistent LUNs flag is set. When set, any masking operation involving this host/host group that would result in inconsistent LUN values, will be rejected. When not set, the storage system will attempt to keep LUN values consistent, but will deviate from consistency if LUN conflicts occur during masking operations. ✓ indicates that the feature is set.
   - **Port Flag Overrides** — Flag indicating if any port flags are overridden for the host. ✓ indicates that there are overridden port flags.
   - **Last Update** — Timestamp of the most recent changes to the host.

   The following controls are available:

   - **Add Hosts** — Adding hosts to host groups on page 332
   - **Remove** — Removing hosts from host groups on page 332
   - **Set flags** — See Setting host/host group flags on page 334.
   - **View Details** — Viewing host/host group details on the previous page

Masking views

Creating masking views

The following explains how to mask volumes on storage systems running Enginuity 5874 or higher. For instructions on masking volumes on storage systems running Enginuity 5773 or lower, refer to Masking volumes on page 369.

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 hosts on page 322, Creating port groups on page 350, and Creating storage groups (Enginuity 5874 - 5876) on page 120.

To create a masking view:

1. Select the storage system.
2. Select **Hosts › Masking view** to open the **Masking view** list view.
3. Click **Create Masking View** 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 **Host**.
6. Select the **Port Group**.
7. Select the **Storage Group**.
8. Optional: To manually set the host LUN addresses:
   a. Click **Set Dynamic LUNs** to open the **Set Dynamic LUN Address** dialog box.
   b. Select a volume, and notice the address displayed in the Starting LUN field. To accept this automatically generated address, click **Apply Starting LUN**. To move to the next available, click **Next Available LUN**.
   c. Click **OK** to close the **Set Dynamic LUN Address** dialog box.
9. Click **OK**.

Renaming masking views

1. Select the storage 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

This procedure explains how to delete masking views from the **Masking Views** list view. In eNAS operating environments, you can also perform this operation from **File Masking Views** page (**System › Dashboard › File Dashboard › File Masking Views**).

1. Select the storage system.
2. Select **Hosts › Masking view** to open the **Masking View** list 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 storage system.
2. Do one of the following:
   - Select **Hosts › Masking view** to open the **Masking Views** list view.
   - Select **Storage › VVols Dashboard › PE Masking Views** to open the **PE 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.
   - **Host** — Name of the associated host.
   - **Port Group** — Name of the associated port group.
   - **Storage Group** — Name of the associated storage group.

   The depending on the options chosen, some of the following controls are available:

   - **Create Masking View** — See Creating masking views on the previous page.
   - **Provision PE to Host** — Provisioning protocol endpoints to hosts
   - **View Connections** — See Viewing masking view connections below.
   - **Delete** — See Deleting masking views above.
   - **View Details** — See Viewing masking view details on page 341.

Viewing masking view connections

This procedure explains how to view masking view connections.

This procedure explains how to perform the operation from the **Masking Views** list view. In eNAS operating environments, you can also perform this operation from **File Masking Views** page (**System › Dashboard › File Dashboard › File Masking Views**).

To view masking view connections:

1. Select the storage system.
2. Select **Hosts › Masking Views** to open the **Masking Views** 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:

- **Hosts** — Name of top-level host.
- **Initiator** — WWN or IQN (iSCSI Qualified Name) ID and the alias of the initiator group.
- **Port Group** — Name of top-level port group.
- **Director:Port** — Storage 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** — WWN or IQN (iSCSI Qualified Name) ID of the initiator.
- **Alias** — Alias of the initiator.
- **Director: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 storage 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.
- **Host** — Name of the associated host.
- **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:

- **Delete** — See Deleting masking views on page 339.
- **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.

Initiators

Masking volumes

The following explains how to mask volumes on storage systems running Enginuity 5773 or lower. For instructions on masking volumes on storage systems running Enginuity 5874 or higher, refer to See Creating masking views on page 338..

You can perform this operation at the initiator or host alias level.

To mask volumes:

1. Select the storage 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 851 and Previewing on page 850.
   - 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 storage systems running Enginuity 5773 or lower.

You can perform this operation at the initiator or host alias level.

1. Select the storage 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 storage 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 storage 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 storage 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**.

   On storage systems running Enginuity 5874 - 5876, node and port names cannot exceed 16 characters.

   On storage systems running HYPERMAX OS 5977 or higher, node and port names cannot exceed 32 characters.
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 storage 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 storage 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 storage system 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 storage 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 storage system and on how you arrived at this view.

- **Initiator** — WWN or IQN (iSCSI Qualified Name) ID of the initiator.
- **Dir:Port** — Storage 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 370
- **Remove** — Removing initiators from host aliases on page 370
- **Mask** — Masking volumes on page 369
- **Set Flags** — See Setting initiator port flags on page 344.
- **Set Attributes** — See Setting initiator attributes on page 344.
- **Rename Alias** — See Renaming initiator aliases on page 345.
- **Replace Initiator** — See Replacing initiators on the previous page.
- **View Details** — See Viewing initiator details below.
- **Remove Masking Entry** — See Removing masking entries on the previous page.

Viewing initiator details

1. Select the storage 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 storage system and on how you arrived at this view.

- **Initiator** — WWN or IQN (iSCSI Qualified Name) ID of the initiator.
- **Dir:Port** — Storage 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** — Wether 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** —Wether 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 369

- **Set Flags** — See Setting initiator port flags on page 344.
- **Set Attributes** — See Setting initiator attributes on page 344.
- **Rename Alias** — See Renaming initiator aliases on page 345.
- **Replace Initiator** — See Replacing initiators on page 346.
- **Removing Masking Entry** — See Removing masking entries on page 346.

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.

Port groups

Creating port groups

Before you begin:

Note the following recommendations:

- Port groups should contain four or more ports.
- Each port in a port group should be on a different director.
- A port can belong to more than one port group. However, for storage systems running HYPERMAX OS 5977 or higher, you cannot mix different types of ports (physical FC ports, virtual ports, and iSCSI virtual ports) within a single port group.

Creating port groups:

1. Select the VMAX system.
2. Select **Hosts › Port Groups** to open the **Port Groups** list view.
3. Click **Create Port Group** 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 storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Port group names are case-insensitive.
5. Select the appropriate filter to filter the port list by iSCSI or FC.
6. 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** — Storage system director and port in the port group.
   - **Identifier** — Port identifier
   - **Port Groups** — 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 flag is set for the port.
7. Click **OK**.

Deleting port groups

1. Select the storage 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 groups

Before you begin:

Note the following recommendations:

- Port groups should contain four or more ports.
- Each port in a port group should be on a different director.
- A port can belong to more than one port group. However, for storage systems running HYPERMAX OS 5977 or higher, you cannot mix different types of ports (physical FC ports, virtual ports, and iSCSI virtual ports) within a single port group.

Adding ports to port groups:

1. Select the storage 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 Port** to open the **Add Ports** dialog box.

   If the port group already contain FC ports, the dialog is populated with all available FC ports. If the port group already contain iSCSI ports, this dialog is populated with all available iSCSI ports.If there are no ports in the port group, select the appropriate filter to filter the port list by iSCSI or FC.
6. Select the available ports from the **Ports to add** list, and click **Add Ports** to add them to the **Ports to Add** list.

   The following properties display:

   - **Dir:Port** — Storage system director and port in the port group.
   - **Identifier**— IQN of an iSCSI target or WWN of an FC port.
   - **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 groups

Before you begin:

Note the following recommendations:

- Port groups should contain four or more ports.
- Each port in a port group should be on a different director.

Removing ports from port groups:

1. Select the storage 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*: You can optionally select to **Unmap** any affected volumes from their respective ports.
8. Click **OK**.

Renaming port groups

1. Select the storage 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 storage 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 storage 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 Port Group** — See Creating port groups on page 350.
   - **View Details** — See Viewing port groups details below.
   - **Delete** — See Deleting port groups on page 351.

Viewing port groups details

1. Select the storage 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 350.
- **Delete** — See Deleting port groups on page 351.
- **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 storage 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** — Storage system director and port in the port group.
   - **Identifier**— IQN of an iSCSI target or WWN of an FC port.
   - **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 Ports** — See Adding ports to port groups on page 351.
   - **Remove** — See Removing ports from port groups on page 352.
   - **View Details** — See Viewing port details on the next page.

Viewing port details

1. Select the storage 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 panels, Properties and Related Objects.

Properties panel

The following properties display:

- **Dir:Port** — Storage system director and port in the port group.
- **Identifier**— IQN of an iSCSI target or WWN of an FC port.
- **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 of 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.
- **Number of IP Interfaces**— Number of IP interfaces associated with the iSCSI target.
- **Number of iSCSI Ports**— Number of physical iSCSI ports associated with IP interfaces which are in turn attached to the iSCSI target.

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.

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 Symmetrix Storage in VMware vSphere Environments*.

To add a new virtual server:

1. Select the storage system or **All Symmetrix** systems.
2. Select **Hosts › Virtual Servers** to open the **Virtual Servers** list view.
3. Click **Add VM Server** 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 **Username** 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 storage 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 storage 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 storage 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 storage 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 storage 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 357.
   - **View Details** — See Viewing the details of a virtual server below.
   - **Change Password** — See Changing the password on a virtual server on the previous page.
   - **Remove** — See Removing a virtual server on page 358.

Viewing the details of a virtual server

1. Select the storage 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 the previous page.
- **Remove** — See Removing a virtual server on page 358.

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 storage 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 the **Base Address** to be assigned to the first volume in the mapping request. Base addresses increases incrementally by one for each volume in the range of volumes being mapped. To view base addresses already in use, click **Show**.
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 851 and Previewing on page 850.
   - 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 storage 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

Hint **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 851 and Previewing on page 850.
   - 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 storage system running Enginuity 56xx or lower. For instructions on assigning aliases on storage systems running 5771 or higher, refer to Assigning an alias range.

Before you begin:

PAV alias must be enabled on the storage system.

To assign an alias address:

1. Select a storage 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 851 and Previewing on page 850.
- 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 storage system running Enginuity 56xx or earlier. For instructions on removing aliases on storage systems running 5771 or higher, refer to See Assigning alias ranges below.

1. Select a storage 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 851 and Previewing on page 850.
   - 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 storage system running 5771 or higher. For instructions on assigning aliases on storage systems running Enginuity 56xx or lower, refer to See Assigning alias addresses on the previous page.

Before you begin:

PAV alias must be enabled on the storage system.

To assign an alias address:

1. Select a storage 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 **Start**ing **Alias** address.
5. Type the **End**ing **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 851 and

Previewing on page 850.

- 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 storage system running 5771 or higher. For instructions on removing aliases on storage systems running Enginuity 56xx or earlier, refer to See Removing alias addresses on the previous page.

Before you begin:

PAV alias must be enabled on the storage system.

To remove alias ranges:

1. Select a storage 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 **Start**ing **Alias** address.
5. Type the **End**ing **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 851 and Previewing on page 850.
   - 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 storage systems running Enginuity 57xx or higher:

1. Select a storage 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 851 and Previewing on page 850.
   - 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 storage systems running Enginuity 57xx or higher:

1. Select a storage 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing CU images

1. Select the storage 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:

- **z/OS Map** — See Mapping CKD volumes on page 362.
- **z/OS Unmap** — See Unmapping CKD volumes on page 362.
- **View Details** — See Viewing CU image details below.
- **Assign Alias Range** — See Assigning alias ranges on page 364.
- **Remove Alias Range** — See Removing alias ranges on page 365.
- **Assign Alias Count** — See Assigning alias counts on page 365.
- **Remove Alias Count** — See Removing alias counts on the previous page.

Viewing CU image details

1. Select the storage 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 Map** — See Mapping CKD volumes on page 362.
- **z/OS Unmap** — See Unmapping CKD volumes on page 362.
- **Assign Alias Range** — See Assigning alias ranges on page 364.
- **Remove Alias Range** — See Removing alias ranges on page 365.
- **Assign Alias Count** — See Assigning alias counts on page 365.
- **Remove Alias Count** — See Removing alias counts on the previous page.

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.

Host aliases

Masking volumes

The following explains how to mask volumes on storage systems running Enginuity 5773 or lower. For instructions on masking volumes on storage systems running Enginuity 5874 or higher, refer to See Creating masking views on page 338..

You can perform this operation at the initiator or host alias level.

To mask volumes:

1. Select the storage 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 851 and Previewing on page 850.
   - 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 369.
   - **Create** — See Creating host aliases on the previous page.
   - **Set Flags** — See Setting initiator port flags on page 344.
   - **View Details** — See Viewing host alias details on the next page.
   - **Delete Host Alias** — See Deleting host aliases above.

Viewing host alias details

1. Select the storage 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 369.
- **Set Flags** — See Setting initiator port flags on page 344.
- **Delete Host Alias** — See Deleting host aliases on the 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.

Host Cache Adapters

Viewing host cache adapters

1. Select the storage system.
2. Select **Hosts › XtremSW Cache Adapters** to open the **XtremSW 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.

CHAPTER 6
Data Protection

This chapter contains the following:

Monitoring protection 376
Local Replication 379
Device Groups 433
Remote Replication 441
Understanding Virtual LUN Migration 484
Replication Groups and Pools 489
Understanding RecoverPoint 517
Open Replicator 541
Understanding Federated Live Migration 552

Monitoring protection

The **Protection Dashboard** provides you with a single place from which to monitor the replication features running on or associated with local storage systems.

Before you begin:

- To access this dashboard, you must be a Monitor, StorageAdmin or higher.
- This dashboard is only available for storage systems running HYPERMAX OS 5977 or higher.
- The information displayed in this dashboard is static, in that it reflects the system's status at the moment you opened the dashboard. To refresh the data, click refresh in the status bar.

To access the Protection Dashboard:

1. Select the storage system.
2. Select **Data Protection › Protection Dashboard**.

Protection Dashboard:

The **Protection Dashboard** includes the following panels:

Summary panel:

- **Remote Arrays** — Number of SRDF connected storage systems.
- **Technologies** — Lists the replication technologies employed on the storage system.

Storage Group Protection panel:

Displays high-level information obtained from protected and unprotected storage groups on the storage system. All of the storage groups on the storage system are organized into the following categories:

Note Partially protected storage groups will not appear on the dashboard. For example, after adding a volume to a storage group protected by ProtectPoint, the storage group is considered partially protected, in which case it will no longer appear in the dashboard. The storage group will remain in this state until you expand it using ProtectPoint, after which it will be considered fully protected, and reappear on the dashboard.

- **Total** — Total number of storage groups on the storage system, regardless of protection. For more information, refer to Viewing total storage groups on page 163.
- **High Availability**— Number of SRDF/Metro storage groups protected by SRDF/Metro. Both R1 and R2 are included in the totals.
- **Remote Replication** — Number of storage groups protected by SRDF. Storage groups involved in multiple SRDF sessions are only counted once. Only R1s or R2s that are an R1 of an R21 setup are displayed. For more information, refer to Viewing remote replication storage groups on page 162.

- **Point in Time** — Number of storage groups protected by TimeFinder. Depending on the storage operating environment, this can include TimeFinder/SnapVX, TimeFinder/Clone, or TimeFinder/Mirror snapshots. Storage groups involved in multiple Timefinder sessions are only counted once. Also, ProtectPoint storage groups with SnapVX snapshots as a source are also included. For more information, refer to Viewing point-in-time storage groups on page 160.
- **Backup** — Number of storage groups with SnapVX snapshots linked to ProtectPoint storage groups. For more information, refer to Viewing backup storage groups (HYPERMAX OS 5977 or higher) on page 159.
- **Unprotected** — Number of storage groups not covered by Continuous, Point in Time or High Availability. For more information, refer to Viewing unprotected storage groups on page 164.

Application Protection Panel:

Allows you to view and manage storage groups that serve as sources in replications sessions.

The following properties display:

- **State** — Replication state severity of the storage group. The value is calculated at the volume level, with the highest (worst) severity being reported for the group. Possible values are:

  Normal — All volumes are in a normal state.

  Warning — One or more volumes in the group are in a warning state, with none higher.

  Error — One or more volumes in the group are in an Error state.

  Replication state severities maps states of the various replication technologies to severities.
- **Storage Group** — Storage group name.
- **Replication Management Group** — Replication management group associated with the storage group. For example, SRDF Group xx (Remote SymmetrixID) for SRDF , the device group name for Clone, or Mirror, and snapshot name for Snap/VX.
- **Technology** — Replication technology used to protect the storage group.

The following controls available, depending on the replication technology protecting the selected storage group:

- Create Snapshot — Creating snapshots on page 409
- Link — Linking to snapshots on page 410
- Restore — Restoring snapshots on page 413
- Unlink — Unlinking from snapshots on page 413
- Manage Protection — Viewing snapshots on page 417
- Establish —Establishing SRDF pairs on page 461
- Split — Splitting SRDF pairs on page 477

- Suspend — Suspending SRDF pairs on page 478
- Restore — Restoring SRDF pairs on page 474
- Resume — Resuming SRDF links on page 470
- Failover — Failing over on page 462
- Failback — Failing back on page 463
- **Manage Protection**— Select a Storage Group and click **Manage Protection**— to navigate to the group list view of the selected technology

Local Replication

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.

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

An overview of a typical clone session is:

1. Create a device group, or add volumes to an existing device group.
2. Create the session; restore the session.
3. Activate the session.
4. View the session's progress.
5. Terminate the session.

For more information on TimeFinder/Clone concepts, refer to the *Solutions Enabler TimeFinder Family CLI Product Guide* and the *TimeFinder Family Product Guide*.

Managing TimeFinder/Clone sessions

The **TimeFinder/Clone** dashboard provides you with a single place to monitor and manage TimeFinder/Clone sessions on a storage system.

Before you begin:

TimeFinder/Clone requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher. On HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents using Clone emulation.

Managing TimeFinder/Clone sessions:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.

   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 Pairs** —See Creating clone copy sessions below.
- **Activate** —See Activating clone copy sessions on page 382.
- **Recreate** —See Recreating clone copy sessions on page 383.
- **View Storage Groups** — Viewing storage groups (Enginuity 5874 - 5876) on page 155
- **View Details** —See Viewing clone copy sessions on page 391.
- **Split** —See Splitting clone volume pairs on page 388.
- **Restore** —See Restoring data from target volumes on page 387.
- **Create Snapshot** — Creating snapshots on page 409
- **Terminate** —See Terminating clone copy sessions on page 389.
- **Set Mode** —See Modifying clone copy sessions on page 385.

Creating clone copy sessions

This procedure explains how to create clone copy sessions.

Before you begin:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- 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_GRPS option in the SYMAPI options file. For more information on enabling SYMAPI options, refer to the *Solutions Enabler CLI Command Reference*.
- Data Domain volumes are not supported.
- The clone copy does not become available to the host until the session is activated.

To create clone copy sessions:

1. Select the storage system.

2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. 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 Pairs** to open the **Create Pairs - Clone Pair** 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 Pairs** to open the **Create Pairs** dialog box.
   c. Click **Set Pairs** to open the **Set TimeFinder Clone 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 Pairs - Clone Pair** dialog box.
4. *Optional*: Select **Use TGT Volumes** to use volumes from a local target list as targets.
5. *Optional*: If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to set the advanced options, as described next.

Setting Advanced options:

1. If performing this operation at the group level, you can optionally select a **Pairing Type** by expanding the menu and selecting 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 Background Copy** — Specifies to start copying tracks in the background at the same time as target I/Os are occurring.
   - **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. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

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:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- You can only activate clone sessions that are in the Created or Recreated state.
- This procedure explains how to perform this operation from the TimeFinder/Clone dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To activate clone copy sessions:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.

3. 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 - Device Group**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 Clone Pair** dialog box.

4. Optional: To attach **Session Options** to the operation, click **Show Advanced** and select any number of options.

5. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

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:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- 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 Enginuity 5876.159.102 or 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 storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. 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 **Recreate** to open the **Recreate - Device Group** 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 **Recreate** to open the **Recreate - Clone Pair**dialog box.

4. *Optional*: To attach **Session Options** to the operation, click **Show Advanced**, and select any number of options.
5. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Creating clone snapshots

This procedure explains how to create and immediately activate clone snapshots.

Before you begin:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- The create 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 create operation.

To create clone snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. 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 **Create Snapshot** to open the **Create Snapshot - Device Group** 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 **Create Snapshot** to open the **Create Snapshot - Clone Pair** dialog box.
4. Specify whether to perform an **Incremental** or **Full** create.
5. If performing this operation at the pair level and performing a Full create, 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 close the dialog box.
6. Select the **Source Type**.
7. Select the **Target Type**.

8. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To specify **Full Snapshot Options**, expand the menu, and do the following:
   a. Select one of the following options. If you are not using the **Full Snapshot 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 storage 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. Do one of the follwing:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Modifying clone copy sessions

This procedure explains how to modify the mode in which a clone copy session is operating.

Before you begin:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- 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 storage system.

2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   Select a group, click more >> , and select **Set Mode** to open the **Set Mode - Device Group** 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 - Clone Pair** dialog box.
4. Select a **Copy Mode**:
   - **Use 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.
   - **Use 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.
   - **Use 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.
5. *Optional*: To set session options, click **Show Advanced**, and select any number of options.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

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:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- 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 5875.139.93, you can perform an incremental restore on volume pairs in a NoCopy/NoDiff clone session.
- With Enginuity 5876.159.102 or higher, you can perform an incremental restore of clone targets to source volumes with active snap and VP snap sessions.
- For a clone session in the Created state, the target volume must be in a fully copied state.

To restore data from a target volume:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   Select a group, click more [>>], and select **Restore** to open the **Restore - Device Group** 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 - Clone Pair** dialog box.

4. 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.
5. If performing the operation at the group level, select the type of source volumes (**Source Type**) and the type of target volumes **Target Type** .
6. To attach **Session Options** to the operation, click **Show Advanced**, and select any number of options.
7. Do one of the following:
   - 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 851 and Previewing on page 850.
   - 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:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- The clone session must be in the Restored state.

To split clone volume pairs:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   Select a group, click more >> , and select **Split** to open the **Split - Device Group**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 - Clone Pair** dialog box.

4. To attach **Session Options** to the operation, click **Show Advanced**, and select any number of options.
5. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Terminating clone copy sessions

This procedure explains how to terminate a clone copy session, thereby deleting the pairing information from the storage system, and removing any hold on the target volume.

Before you begin:

- TimeFinder/Clone requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- You need a clone copy session in any pair state.
- Terminating a session while the pairs are in the CopyOnAccess, CopyOnWrite, or CopyInProg 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 storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder/Clone** dashboard.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   Select a group, click more >> , and select **Terminate** to open the **Terminate - Device Group** 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** to open the **Terminate - Clone Pair** dialog box.
4. If performing the operation at the group level, select the type of source volumes (**Source Type**) and the type of target volumes **Target Type** .
5. To attach **Session Options** to the operation, click **Show Advanced**, and select any number of options.
6. Do one of the following:
   - 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 851 and Previewing on page 850.

- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing clone copy sessions

1. Select the storage system.
2. Select **Data Protection › TimeFinder › Time/Finder/Clone** to open the **TimeFinder/Clone** view.
3. 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.
     - **Timestamp** — Date and time the snapshot was created.

The following controls are available:

- **Create Pairs** —See Creating clone copy sessions on page 380.
- **Activate** —See Activating clone copy sessions on page 382.
- **Recreate** —See Recreating clone copy sessions on page 383.
- **View Details** — Viewing clone session details on the next page
- **Split** —See Splitting clone volume pairs on page 388.
- **Restore** —See Restoring data from target volumes on page 387.
- **Create Snapshot** —See Creating clone copy sessions on page 380.
- **Terminate** —See Terminating clone copy sessions on page 389.
- **Set Mode** —See Modifying clone copy sessions on page 385.

Viewing clone copy sessions

1. Select the storage system.
2. Select **Data Protection › TimeFinder › Time/Finder/Clone** to open the **TimeFinder/Clone** view.
3. 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.
  - **Timestamp** — Date and time the snapshot was created.

The following controls are available:

- **Create Pairs** —See Creating clone copy sessions on page 380.
- **Activate** —See Activating clone copy sessions on page 382.
- **Recreate** —See Recreating clone copy sessions on page 383.
- **View Details** — Viewing clone session details below
- **Split** —See Splitting clone volume pairs on page 388.
- **Restore** —See Restoring data from target volumes on page 387.
- **Create Snapshot** —See Creating clone copy sessions on page 380.
- **Terminate** —See Terminating clone copy sessions on page 389.
- **Set Mode** —See Modifying clone copy sessions on page 385.

Viewing clone session details

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Clone** to open the **TimeFinder Clone** view.
3. Select a group, and click **View Details** to open its **Clone pair** list view.
4. Select a pair and click **View Details** to open its **Details** view.
5. Use the clone session **Details** view to view and manage a clone session.

   The following properties display:

- **Group** — The type of group. Possible values are: Regular, R1, R2, or R21.
- **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.
- **Pair State** — The session state of the pair.
- **Source protected tracks** — Number of protected tracks on the source volume.
- **Source modified tracks** — Number of tracks modified on the source volume.
- **Target protected tracks** — Number of protected tracks on the target volume.
- **Target modified tracks** — Number of modified tracks on the target volume.
- **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.
- **Timestamp** — Date and time the pair was created.

Clone copy session options

The following table describes the TimeFinder/Clone session options:

Table 5: 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
Concurrent	Performs the action for an additional clone pair in a group.	Create Recreate Establish

Session option	Description	Available with action
		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 storage 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 storage system.	Create Establish
Precopy	Copies tracks in the background before the clone session is activated. Used with the create and recreate actions.	Create 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

Session option	Description	Available with action
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 storage volumes that are not host-accessible and can only be accessed through the virtual volumes that point to them. SAVE volumes provide pooled physical storage for virtual volumes.

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

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

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

The following are the basic actions 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 *Solutions Enabler TimeFinder Family CLI Product Guide* and the *TimeFinder Family Product Guide*.

Managing TimeFinder/Snap sessions

The **TimeFinder/Snap** dashboard provides you with a single place to monitor and manage TimeFinder/Snap sessions on a storage system.

Before you begin:

TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.

Managing TimeFinder/Snap sessions:

1. Select the Symmetrix system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.

   The following properties display:

   - **Source Group** — Groups containing volumes using TimeFinder/Snap.
   - **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 session state of the pair.
   - **Group Type** — The type of group. Possible values are: Regular, R1, R2, or R21.
   - **Group Valid** — Whether the group is valid or invalid.

   The following controls are available:

   - **Create Pairs** — See Creating virtual copy sessions on the facing page.
   - **Activate** — See Activating virtual copy sessions on page 398.
   - **Terminate** — See Terminating virtual copy sessions on page 404.
   - **View Storage Groups** — Viewing storage groups (Enginuity 5874 - 5876) on page 155
   - **View Details** — Viewing snap pair details
   - **Restore** — See Restoring virtual copy sessions on page 403.

- **Recreate** — See Recreating virtual copy sessions on page 402.
- **Create Snapshot** — Creating snapshots on page 399
- **Duplicate** — See Duplicating virtual copy sessions on page 401.

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:

- TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.
- 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 storage systems running Enginuity 5875- 5876, 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.
- Data Domain volumes are not supported

To create virtual copy sessions:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.

3. 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 Pairs** to open the **Create - Device Group** dialog box.

   b. Select a **Snap Pool**.

   c. Select the type of **Source Volumes**.

   d. Select the type of **Target Volumes**.

   Pair level:

   a. Select a group, and click **View Details** to open its **Details** view.

   b. Select one or more pairs, and click **Create Pairs** to open the **Create Pairs** 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.

4. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **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 storage 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. Do one of the follwing:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

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.

Before you begin:

TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.

To activate virtual copy sessions:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.
3. 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.
4. Select the type of **Source Volumes** to use.
5. Select the type of **Target Volumes** to use.
6. 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**.

Creating snapshots

This procedure explains how to create and immediately activate virtual copy sessions.

Before you begin:

TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.

Creating snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.

3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   Select a group, click more >> , and select **Create Pairs** to open the **Create Pairs** 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 **Create Pairs** to open the **Create Pairs** dialog box.

4. *Optional*: Select **Use TGT Volumes** to use volumes from a local target list as targets.
5. *Optional*: If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
6. 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 storage 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:

- TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.
- 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 storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.
3. Select a group, and click **View Details** to open its **Snap Sessions List** view.
4. 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:

- TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.
- 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 storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.

3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   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.

4. *Optional*: Select **Use TGT Volumes** to use volumes from a local target list as targets.
5. *Optional*: If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
6. 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:

- TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.
- For storage systems running Enginuity 5875 or higher, you can use this feature to recreate multivirtual snap sessions from thin and standard volumes.
- This feature can only be used on sessions that have been previously activated.

To recreate virtual copy sessions:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

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.

4. *Optional*: Select **Use TGT Volumes** to use volumes from a local target list as targets.
5. *Optional*: If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
6. 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:

- TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.
- With Enginuity 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 Enginuity 5876.159.102 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 storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:

   Group level:

   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.
4. 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.
5. 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.
6. 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.

Before you begin:

TimeFinder/Snap requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.

To terminate virtual copy sessions:

1. Select the storage system.

2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.
3. 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** to open the **Terminate Sessions** dialog box.
4. *Optional*: Select **Use TGT Volumes** to use volumes from a local target list as targets.
5. *Optional*: If you selected **Use TGT Volumes**, you can also select to **Use BCV Volumes** as the source.
6. 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 storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder/Snap** dashboard.
3. 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** — The logical 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.
- **Timestamp** — Date and time the snapshot was created.

The following controls are available:

- **Create Pairs** — See Creating virtual copy sessions on page 397.
- **Activate** — See Activating virtual copy sessions on page 398.
- **Terminate** — See Terminating virtual copy sessions on page 404.
- **View Details** — Viewing snap pair details
- **Detach** — See Attaching and detaching preferred pairs on page 400.
- **Attach** — See Attaching and detaching preferred pairs on page 400.
- **Duplicate** — See Duplicating virtual copy sessions on page 401.
- **Create Snapshot** — See Creating snapshots on page 399.
- **Recreate** — See Recreating virtual copy sessions on page 402.
- **Restore** — See Restoring virtual copy sessions on page 403.

Viewing snap pair details

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Snap** to open the **TimeFinder Snap** view.
3. Select a group, and click **View Details** to open its **Snap Sessions** list view.
4. Select a pair and click **View Details** to open its **Details** view.
5. 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.

Snap session options

The following table describes the TimeFinder/Snap session options:

Table 6: TimeFinder/Snap 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
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 SnapVX

Managing TimeFinder SnapVX

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

The **TimeFinder/SnapVX** dashboard provides a single place from you can manage TimeFinder SnapVX snapshots and their associated storage groups.

Before you begin:

- The storage system must be running HYPERMAX OS 5977 or higher.
- TimeFinder/SnapVX operations are not supported on *working* ProtectPoint snapshots.
  TimeFinder/SnapVX operations are, however, supported to help repair failing ProtectPoint snapshots.

To access the TimeFinder/SnapVX dashboard:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.

TimeFinder/SnapVX dashboard:

The following properties display:

- **Storage Groups** — Storage group associated with the snapshot.
- **Capacity** — Total capacity of the storage group.
- **Snapshots** — Number of snapshots associated with storage group.
- **Creation Time** — Date/time the most recent snapshot was created.

The following controls are available:

- **Create Snapshot** — Creating snapshots on the facing page
- **Link** — Linking to snapshots on page 410
- **Restore** — Restoring snapshots on page 413
- **View Snapshots** — Viewing snapshots on page 417
- **Relink** — Relinking to snapshots on page 412

- **Unlink** —Unlinking from snapshots on page 413
- **Set Mode** —Setting copy mode for snapshots on page 416
- **Set Days to Live** —Setting snapshots to automatically terminate on page 414
- **Terminate** —Terminating snapshots on page 415

Creating snapshots

This procedure explains how to create TimeFinder SnapVX snapshots.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.
- The maximum number of snapshots per source volume is 256.
- Snapshots off of linked targets are permitted only after the volume is fully defined.
- You can perform this operation from the either of the following dashboards: **TimeFinder/SnapVX**, **Storage**, or **Protection**. Depending on the location from which you are performing this operation, some of the following steps may not apply.

To create snapshots:

1. Select the storage system.
2. Do the following, depending on the location from which you want to perform the procedure:
   TimeFinder/SnapVX dashboard:
   a. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
   b. Select a storage group and click **Create Snapshot** to open the **Create Snapshot** dialog box.

   Storage Groups Dashboard:
   a. Select **Storage › Storage Groups Dashboard** to open the **Storage Groups Dashboard**.
   b. In the **Storage Groups** panel, click **Total** or **Unprotected** to open the corresponding list view.
   c. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.
   d. If not already selected, select **Point In Time Protection**, and then select **Using SnapVX**.
   e. Click **Next**.

   Protection dashboard:
   a. Select **Data Protection › Protection Dashboard** to open the **Protection Dashboard**.
   b. In the **Storage Group Protection** panel, click **Total** to open the **All Storage Groups** list view.
   c. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.

d. If not already selected, select **Point In Time Protection**, and then select **Using SnapVX**.

e. Click **Next**.

1. Select whether to create a new snapshot or reuse an existing snapshot.
2. If reusing an existing snapshot, select it from the list. When using this method, Unisphere for VMAX assigns generation numbers to the snapshots in the order in which they were created (latest = generation 0, previous incrementing by one). This naming convention allows you to differentiate point-in-time copies of the same volumes.

   CAUTION It is the users responsibility to manage the snapshot names being used. If snapshots are being applied to parent and child storage groups individually, care should be taken to never use the same snapshot name at different levels of the storage group construct. The same applies if some of the volumes are in multiple storage groups being snapshotted; the same snapshot names should also be avoided across the different storage groups.

3. Do one of the following:
   - Click **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - Expand **Add to Job List** and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to continue setting the advanced options, as described next.

Setting Advanced options:

1. To configure the snapshot to automatically terminate after a period of time, specify the number of **Days to live**. Once the time has expired, Unisphere for VMAX will automatically terminate the snapshot, provided that it is not linked to any target volumes. If an expired snapshot is linked, Unisphere for VMAX will wait until the last link has been removed before terminating the snapshot. To override this behavior, select the **Force** option, which will allow Unisphere for VMAX to terminate the snapshot regardless of whether it is linked. When **Days to live** is set to 0, the snapshot will remain until you manually terminate it.
2. To force the operation even though one or more volumes may not be in the normal, expected states, select **Force**.
3. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - Expand **Add to Job List** and click **Run Now** to perform the operation now.

Linking to snapshots

This procedure explains how to link one or more host-mapped target volumes to a snapshot, thereby making the snapshot's point-in-time data available to applications running on the host.

Snapshots can be linked to target volumes in the following modes:

- NoCopy mode — Creates a temporary, space-saving snapshot of only the changed data on the snapshot's Storage Resource Pool (SRP). Target volumes linked in this mode will not retain data after the links are removed. This is the default mode. This mode cannot be used when either the source or link target volume is a Data Domain volume.
- Copy mode — Creates a permanent, full-volume copy of the data on the target volume's SRP. Target volumes linked in this mode will retain data after the links are removed.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.
- The targets must not be linked to any other snapshots.
- The target volume must be of equal or greater size than the source volume.
- Any pre-existing data that was exclusive to the target will be lost during a link or relink.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To link to snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group and clink **Link** to open the **Link to Storage Group** dialog box.
4. Select the **Snapshot Name** and **Creation Date** (0 is the latest).
5. Specify whether to link to a new target storage group (one not already linked to a snapshot) or an existing target storage group.
6. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - 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 create a permanent, full-time copy of the data on the target volume's SRP, select **Copy**.
2. To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.
3. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.

- Expand **Add to Job List** and click **Run Now** to perform the operation now.

Relinking to snapshots

This procedure explains how to unlink a target storage group from a snapshot, and then automatically link it to another snapshot. After a relink operation, the copy between the original linked snapshot and the newly linked snapshot is differential.

You can also relink storage group to the same snapshot, thereby refreshing the point-in-time copy on the target storage group when it's been modified by host writes.

Before you begin:

- To perform this operation, you must be a StorageAdmin.

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

- To relink in Copy mode:
  - The original link must be fully copied prior to the relink.
  - The copy will be differential between the original linked snapshot and the newly linked snapshot.
- Any pre-existing data that was exclusive to the target will be lost during a link or relink.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To relink snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group, click more >> and select **Relink** to open the **Relink** dialog box.
4. Select the new **Snapshot Name** and **Generation** (0 is the latest).
5. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - 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 create a permanent, full-time copy of the data on the target volume's SRP, select **Copy**.
2. To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.
3. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on

page 851.

- Expand **Add to Job List** and click **Run Now** to perform the operation now.

Unlinking from snapshots

This procedure explains how to unlink target volumes from their snapshots.

For instructions on unlinking target volumes, and then automatically linking to other snapshots, refer to Relinking to snapshots on the previous page.

Before you begin:

- To perform this operation, you must be a StorageAdmin or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To unlink from snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group, click more >> and select **Unlink** to open the **Unlink** dialog box.
4. Select the **Snapshot Name** and **Creation Date** (0 is the latest).
5. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - 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 force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.
2. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - Expand **Add to Job List** and click **Run Now** to perform the operation now.

Restoring snapshots

This procedure explains how to restore snapshot data back to the original source volumes. TimeFinder SnapVX restore operations are inherently differential, meaning that only tracks that have changed since the snapshot was created are copied back to the source volumes.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To restore snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group and clink **Restore** to open the **Restore** dialog box.
4. Select the **Snapshot Name** and **Creation Date** (0 is the latest).
5. To force the operation even though one or more volumes may not be in the normal, expected states, click **Show Advanced** , and select **Force**.
6. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - 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 snapshots to automatically terminate

This procedure explains how to configure snapshots to automatically terminate after a period of time.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To set snapshots to automatically terminate:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group, click more >> and select **Set Days to Live** to open the **Set Days to Live**dialog box.

4. Select the **Snapshot Name** and **Generation**.
5. Select the number of **Days to live**. Once the time has expired, Unisphere for VMAX will automatically terminate the snapshot, provided that it is not linked to any target volumes. If an expired snapshot is linked, Unisphere for VMAX will wait until the last link has been removed before terminating the snapshot. To override this behavior, select the **Force** option, which will allow Unisphere for VMAX to terminate the snapshot regardless of whether it is linked. To remove the **Days to live** attribute, select **None**.
6. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - 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 force the operation even though one or more volumes may not be in the normal, expected states, select **Force**.
2. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - Expand **Add to Job List** and click **Run Now** to perform the operation now.

Terminating snapshots

This procedure explains how to remove snapshots from a system.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.
- The snapshot must not have any links.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To terminate snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group, click more >> and select **Terminate** to open the **Terminate** dialog box.
4. Select the **Snapshot Name** and **Generation**.
5. Click one of the following:

- **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
- 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 force the operation even though one or more volumes may not be in the normal, expected states, select **Force**.
2. To force the operation when the operation would normally be rejected, select **SymForce**.

   CAUTION Use extreme caution with this option. If used when a link is copy in progress or when a restore is restore in progress, this will cause an incomplete copy and data on the copy target would not be usable.

3. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - Expand **Add to Job List** and click **Run Now** to perform the operation now.

Setting copy mode for snapshots

This procedure explains how to change the copy mode used when linking target volumes and snapshots.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.
- This procedure explains how to perform this operation from the TimeFinder/SnapVX dashboard. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

To set the copy mode for snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select the storage group, click more >> and select **Set Mode** to open the **Set Mode** dialog box.
4. Select the **Snapshot Name** and **Generation**.
5. Select a new mode:
   - **No Copy** — Creates a temporary, space-saving snapshot of only the changed data on the snapshot's Storage Resource Pool (SRP). Target volumes linked in this mode will not retain data after the links are removed. This is the default mode.

- **Copy** — Creates a permanent, full-volume copy of the data on the target volume's SRP. Target volumes linked in this mode will retain data after the links are removed.

6. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - 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 force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.
2. Click one of the following:
   - **Add to Job List** to perform the operation at a later time, as described in Scheduling jobs on page 851.
   - Expand **Add to Job List** and click **Run Now** to perform the operation now.

Viewing snapshots

This procedure explains how to view and manage snapshots of a storage group.

Before you begin:

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

To view snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.

   The storage group **Snapshots** list view allows you to view and manage the snapshots associated with a storage group.

   The following properties display:

   - **Snapshot** — Name of the snapshot.
   - **Creation Time** — Date, time, and generation number for the snapshot.
   - **Linked** — Whether the snapshot is linked to another storage group. A checkmark (✓) indicates that the snapshot is linked.
   - **Restored** — Whether the snapshot is restored to the source. A checkmark (✓) indicates that the snapshot is restored.
   - **Tracks** — Number of tracks that have been allocated to preserve the point in time copy of the source volume since the snapshot was created.

- **Non-Shared Tracks** — Number of tracks uniquely allocated for the snapshot's delta.

The following controls are available:

- **Create Snapshot** —Creating snapshots on page 409
- **Link** —Linking to snapshots on page 410
- **Restore** —Restoring snapshots on page 413
- **View Details** —Viewing snapshot details below
- **Relink** —Relinking to snapshots on page 412
- **Unlink** — Unlinking from snapshots on page 413
- **Set Mode** —Setting copy mode for snapshots on page 416
- **Set Days to Live** —Setting snapshots to automatically terminate on page 414
- **Terminate** —Terminating snapshots on page 415

Viewing snapshot details

Before you begin:

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

To view snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.
4. Select a snapshot and select **View Details** to open the snapshot **Details** view.

   The snapshot **Details** view allow you to view and manage a snapshot. It contains Properties and Related Objects panels.

Properties panel:

The following properties display:

- **Name** — Name of the snapshot. To rename the snapshot, type a new name over the existing and click **Apply**. Snapshot names must be unique from other napshots on the system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Snapshot names are case-sensitive. You are responsible for naming snaphots,and you should avoid using the same snapshot names in different containers like a storage group.
- **Generation** — Generation number assigned to the snapshot. This number is used to differentiate between point-in-time copies of the same name and same volumes. Unisphere for VMAX assigns generation numbers to the snapshots in the order in which they were created (latest = generation 0, previous incrementing by one).

- **Creation Date** — Date and time the snapshot was created.
- **Expiry Date** — Date and time the snapshot is set to automatically terminate. If the snapshot is not set to automatically terminate, this field displays N/A.
- **State** — Snapshot state.

The following controls are available:

- **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 snapshot. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Links - 2** will open a view listing the two links associated with the snapshot.

Viewing snapshot links

This procedure explains how to view storage groups containing the linked volumes.

Before you begin:

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

To view snapshot source volumes:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.
4. Select a snapshot and select **View Details** to open the snapshot **Details** view.
5. In the **Related Objects** panel, click **Links** to open the snapshot **Links** list view.

   The snapshot **Links**list view allow you to view and manage the storage groups containing the linked volumes.

   The following properties display:

   - **Storage Group** — Name of the storage group.
   - **State** — Snapshot state.
   - **Timestamp** — Date and time the snapshot was created.
   - **Tracks** — Number of source tracks that the host has not yet overwritten.
   - **Track Size** — Track size in bytes.

- **Percent Copied** — Percent of the total source volume size that has been overwritten by the host.

The following controls are available:

- **Link** —Linking to snapshots on page 410
- **Relink** —Relinking to snapshots on page 412
- **View Details** — Viewing snapshot link details below

Viewing snapshot link details

This procedure explains how to view snapshot link pairs.

Before you begin:

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

To view snapshot source volumes:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.
4. Select a snapshot and select **View Details** to open the snapshot **Details** view.
5. In the **Related Objects** panel, click **Links** to open the snapshot **Links** list view.
6. Select a link and click **View Details** to open the snapshot link **Details** view.

   The snapshot link **Details** view allow you to view and manage the linked volume pairs.

   The following properties display:

   - **Source Volume** — Name of the source volume.
   - **Link Volume** — Name of the linked volume.
   - **State** — Snapshot state.
   - **Timestamp** — Date and time the snapshot was created.
   - **Flags** — Snapshot flags. Possible values are: Failed, Links, Restored, Defined (FLRD).
   - **Percent Copied** — Percent of the total source volume size that has been overwritten by the host.

   The following controls are available:

   - **Relink** —Relinking to snapshots on page 412
   - **Unlink** — Unlinking from snapshots on page 413

Viewing snapshot source volumes

This procedure explains how to view the source volumes in snapshot.

Before you begin:

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

To view snapshot source volumes:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.
4. Select a snapshot and select **View Details** to open the snapshot **Details** view.
5. In the **Related Objects** panel, click **Source Volumes** to open the snapshot **Source Volumes** list view.

   The snapshot **Source Volumes** view allow you to view and manage the source volumes in a snapshot.

   The following properties display:

   - **Name** — Name of the volume.
   - **State** — Snapshot state.
   - **Creation Date** — Date and time the snapshot was created.
   - **Flags** — Snapshot flags. Possible values are: Failed, Links, Restored, Defined (FLRD)
   - **Percent Copied** — Percent of the total source volume size that has been overwritten by the host.

   The following controls are available:

   - **Link** —Linking to snapshots on page 410
   - **Restore** —Restoring snapshots on page 413
   - **Relink** —Relinking to snapshots on page 412
   - **View Details** — Viewing snapshot source volume details below
   - **Unlink** — Unlinking from snapshots on page 413
   - **Set Mode** —Setting copy mode for snapshots on page 416
   - **Set Days to Live** —Setting snapshots to automatically terminate on page 414
   - **Terminate** —Terminating snapshots on page 415

Viewing snapshot source volume details

This procedure explains how to view the details on a snapshot source volume.

Before you begin:

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

To view snapshot source volumes:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.
4. Select a snapshot and select **View Details** to open the snapshot **Details** view.
5. In the **Related Objects** panel, click **Source Volumes** to open the snapshot **Source Volumes** list view.
6. Select the volume and click **View Details** to open the snapshot source volume **Details** view.

   The snapshot source volume **Details** view allow you to view and manage the source volume in a snapshot.

   The following properties display:

   - **Name** — Name of the volume.
   - **State** — Snapshot state.
   - **Timestamp** — Date and time the snapshot was created.
   - **Child Storage Group** — Name of the child storage group containing this volume. If the volume is not in a child storage group, this field will display N/A.
   - **Flags** — Snapshot flags. Possible values are: Failed, Links, Restored, Defined (FLRD)
   - **Tracks** — Number of source tracks that the host has not yet overwritten.
   - **Track Size** — Track size in bytes.

   The following controls are available:

   - **Link** —Linking to snapshots on page 410
   - **Restore** —Restoring snapshots on page 413
   - **Relink** —Relinking to snapshots on page 412
   - **Unlink** — Unlinking from snapshots on page 413
   - **Set Mode** —Setting copy mode for snapshots on page 416
   - **Set Days to Live** —Setting snapshots to automatically terminate on page 414
   - **Terminate** —Terminating snapshots on page 415

Viewing snapshot source volume linked volumes

This procedure explains how to view linked volumes for a snapshot source volume.

Before you begin:

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

To view snapshot source volumes:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/SnapVX** to open the **TimeFinder/SnapVX** dashboard.
3. Select a storage group and select **View Snapshots** to open the storage group **Snapshots** list view.
4. Select a snapshot and select **View Details** to open the snapshot **Details** view.
5. In the **Related Objects** panel, click **Source Volumes** to open the snapshot **Source Volumes** list view.
6. Select the volume and click **View Details** to open the snapshot source volume **Details** view.
7. In the **Related Objects** panel, click **Link Volumes** to open the snapshot source volume **Link Volumes** list view.

   The snapshot source volume **Link Volumes** list view allow you to view and manage the linked volumes for a snapshot source volume.

   The following properties display:

   - **Name** — Name of the volume.
   - **Storage Group** — Storage group that contains the target volume.
   - **State** — Snapshot state.
   - **Creation Time** — Date and time the snapshot was created.
   - **Flags** — Snapshot flags. Possible values are: Failed, Links, Restored, Defined (FLRD).
   - **Percent Copied** — Percent of the total source volume size that has been overwritten by the host.

   The following controls are available:

   - **Link** —Linking to snapshots on page 410
   - **Unlink** — Unlinking from snapshots on page 413

TimeFinder/Mirror

Managing TimeFinder/Mirror sessions

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

Before you begin:

- TimeFinder/Miorror requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.
- TimeFinder operations are not supported on ORS control volumes on storage systems running HYPERMAX OS 5977 or higher.

Managing TimeFinder/Mirror sessions:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** dashboard.

   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:

   - **Create Snapshot**— See Creating Snapshots below.
   - **Restore** — See Restoring BCV pairs on the facing page.
   - **View Storage Groups** — Viewing storage groups (Enginuity 5874 - 5876) on page 155
   - **Cancel** — See Cancelling BCV pairs on page 427.
   - **Split** — See Splitting BCV pairs on page 427.
   - **View Details** — See Viewing mirror pair details on page 429.

Creating Snapshots

Before you begin:

- TimeFinder/Miorror requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.
- Data Domain volumes are not supported.

To create snapshots:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** dashboard.
3. Do the following, depending on whether you want to perform this operation at the group level or the pair level:

   Group level:

   Select a device group, and click **Create Snapshot** to open the **Create Snapshot - 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 **Create Snapshot** to open the **Create Snapshot - Mirror Pair** dialog box.
4. Select a **Snapshot 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.
5. 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 **Create Snapshot - Mirror Pair** dialog box.
6. 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 Snapshot Options**, expand the menu, and select one of the following options. If you are not using the **Full Snapshot 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 storage 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. For option descriptions, refer to See TimeFinder/Mirror session options on page 430.
3. Click **OK**.

Restoring BCV pairs

This procedure explains how to copy data from the BCV volumes to the standard volumes.

Before you begin:

- TimeFinder/Miorror requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.

To restore BCV pairs:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** dashboard.
3. Do the following, depending on whether you want to perform this operation at the group level or the pair level:

   Group level:

   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.
4. 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.
5. 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.
6. 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 storage 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 430.

3. Click **OK**.

Splitting BCV pairs

This procedure explains how to split paired volumes to where each holds separate valid copies of the data.

Before you begin:

TimeFinder/Miorror requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.

To split BCV pairs:

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** dashboard.
3. 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.
4. 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 430.
2. Click **OK**.

Cancelling BCV pairs

This procedure explains how to cancel the relationship between volumes in a BCV pair.

Before you begin:

TimeFinder/Miorror requires Enginuity version 5773 - 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.

To cancel BCV pairs:

1. Select the storage system.

2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** dashboard.
3. Do the following, depending on whether you want to perform this operation at the group level or the pair level:

   Group level:

   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.
4. 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 430.
2. Click **OK**.

Viewing mirror pairs

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** view.
3. Select a device group and click **View Details** to open the **Mirror Pairs** list view.
4. Use the **Mirror Pairs** list view to display and manage TimeFinder/Mirror sessions.

   The following properties display:

   - **Show Group Details** — Lists additional group details, including:
     - **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.
  - **Percent Copied** — The percentage of copying that is complete.
  - **Timestamp** — Date and time the pair was established.

The following controls are available:

- **Create Snapshot** — See Creating Snapshots on page 424.
- **Restore** — See Restoring BCV pairs on page 425.
- **Split** — See Splitting BCV pairs on page 427.
- **View Details** —Viewing mirror pair details below
- **Cancel** — See Cancelling BCV pairs on page 427.

Viewing mirror pair details

1. Select the storage system.
2. Select **Data Protection › TimeFinder › TimeFinder/Mirror** to open the **TimeFinder/Mirror** view.
3. Select a device group and click **View Details** to open the **Mirror Pairs** list view.
4. Select a pair and click **View Details** to open its **Details** view.
   - **Group** — Group name.
   - **Source Volume** — Hexadecimal ID of the source volume.
   - **Source LDev** — Logical name of the source volume.
   - **Target Volume** — Hexadecimal ID of the target volume.
   - **Target LDev** — Logical name of the target volume.
   - **Pair State** — 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 pre-copy operation has completed one cycle.

. = The pre-copy operation has not completed one cycle.

- **Number of source invalid tracks** — Number of invalid tracks owed to the target volume.
- **Number of target invalid tracks** — Number of invalid tracks owed to the source volume.
- **Percent Copied** — Percentage of copying complete.

TimeFinder/Mirror session options

The following table describes the TimeFinder/Mirror session options:

Table 7: TimeFinder/Mirror session options

Session option	Description	Available with action
Bypass	Bypasses the storage 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
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	Create Establish Activate

Session option	Description	Available with action
	caution when checking this option because improper use may result in data loss.	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 storage system 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 within the same RDF (RA) group on a Symmetrix system.	Full Establish
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	Full Establish Incremental Establish Split Full Restore Incremental Restore

Session option	Description	Available with action
	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.	
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 checking this option because improper use may result in data loss.	Full Establish Incremental Establish Split Full Restore Incremental Restore

Device Groups

Creating device groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.
4. Click **Create** to open the **Create Device Group** wizard.
5. Type a **Device Group Name.**
6. Select a **Device Group Type**. Possible values are:
   - **Regular** — Group can only contain REGULAR volumes.
   - **R1** — Group can only contain R1 volumes.
   - **R2** — Group can only contain R2 volumes.
   - **R21** — Group can only contain R21 volumes.
   - **Any** — Group can contain any volume type.
7. Click **Next.**
8. Select the **Source** of the volumes to use when creating the group; either manual selection, or all the volumes in a storage group.
9. 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**.
10. Click **Next**.
11. Select how to specify the **Target Volumes**, either manually or automatically.

12. Do the following, depending on how you are specifying the target volumes:

    Automatically:

    a. *Optional*: Select to replicate the source volumes using **TimeFinder/Snap**, **TimeFinder Mirror**, or **TimeFinder/Clone**. The required devices (if they are not found to be already existing and unused) will be created. The BCV devices will be automatically created for the TimeFinder Mirror device group. The VDEV devices will be automatically created for the TimeFinder/Snap device group. The required devices will be automatically created for the TimeFinder/Clone device group.

    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**.

13. Click **Next**.
14. 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.
15. Click **Finish**.

A window appears that displays the progress of the wizard's tasks.

Adding volumes to device groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.
4. Select the device group and click **Edit** to open the **Add Volumes to Device Group** view.
5. From the list of available volumes select the volumes and click **Add to Group.**
6. Click **OK**.

Removing volumes from device groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.
4. Select the device group and click **View Details** to open the **Details** view.
5. In the **Related Objects** panel, click a volume type to open the list view for the volume type.
6. Select one or more volumes and click **Remove**.
7. Click **OK** in the confirmation message.

Enabling consistency protection

This procedure explains how to enable consistency protection.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To enable consistency:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
4. Select a group, click more >> , and select **Enable Consistency** to open the dialog box.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Disabling consistency protection

This procedure explains how to disable consistency protection.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To disable consistency:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
4. Select a group, click more >> , and select **Disable Consistency** to open the dialog box.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the

advanced options and click **OK**.

Renaming device groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.
4. Select the device group from the list and click **View Details** to open the **Details** view.
5. In the **Name** field, enter the new device group name and click **Apply** to save changes.

Deleting device groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.
4. Select the device group and click **Delete**.
5. Click **Delete** in the confirmation message.

Viewing device groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.

   Use the **Device Group** list view to view and manage device groups.

   The following properties display, depending on the operating environment:

   - **Name** — User-defined device group name.
   - **Group Type** — Device configuration of the devices in the group. Possible values are: Regular, R1, R2, R21, or Any.
   - **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.
   - **Gatekeepers** — Number of gatekeeper devices in the device group. (Does not apply/display with HYPERMAX OS 5977.)
   - **Valid** — Indicates whether the device group is valid.

   The following controls are available, depending on the operating environment:

   - **Create** — See Creating device groups on page 433.
   - **Edit** — See Adding volumes to device groups on page 434.
   - **Delete** — See Deleting device groups on the previous page.
   - **View Details** — See Viewing device group details below.
   - **Replication QoS** — Setting replication QoS on page 232
   - **Assign Dynamic Cache Partition** — Assigning dynamic cache partitions on page 875 (Does not apply/display with HYPERMAX OS 5977.)
   - **Assign Symmetrix Priority** — Assigning Symmetrix priority on page 223
   - **Set Optimized Read Miss** — Setting optimized read miss on page 227 (Does not apply/display with HYPERMAX OS 5977.)

Viewing device group details

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools**.
3. Click **Device Groups** to open the **Device Group** list view.
4. 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, depending on the operating environment:

- **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** — Storage system serial number ID.
- **Remote Symmetrix ID** — storage system serial number ID of the remote storage system.
- **RBCV Symmetrix ID** — storage system serial number ID for the remote BCV.
- **RRBCV Symmetrix ID** — storage system serial number ID for remote RBCV.
- **Number of Associated Gatekeepers** — Number of gatekeeper devices in the device group.
- **Number of STD Volumes in Group** — Number of standard devices in the device group.
- **Number of Locally-Associated BCVs** — Number of local 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 TGTs** — Number of local target volumes associated with the device group.
- **Number of Remotely-Associated BCVs (STD SRDF)** — Number of remote BCV devices associated with the device group.
- **Number of Remotely-Associated BCVs (BCV SRDF)** — Number of BCV devices, associated with the device group, to be paired with remotely-attached BCV devices.
- **Number of Remotely-Associated RBCVs (RBCV SRDF)** — 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.
- **Pacing Capable** — Indicates if the device group allows write pacing capability.
- **Group-level Pacing State**— Indicates if the device group is write pacing enabled or disabled.
- **Volume-level Pacing State** — Indicates if the volumes in the device group are write pacing enabled or disabled.
- **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 Write Pacing Exempt Volumes** — Indicates if the volumes in the device group have write pacing exemption capability enabled or disabled.
- **Cacheless Read Miss** — Cacheless read miss status. Possible values are:
  - **Off** — Feature is disabled.
  - **System Default** — storage system determines the appropriate optimized read miss mode.

The following controls are available, depending on the operating environment:

- **Create** — See Creating device groups on page 433.
- **Edit** — See Adding volumes to device groups on page 434.
- **Delete** — See Deleting device groups on page 436.
- **Replication QOS** — Setting replication QoS on page 232
- **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 will display with inactive links if the selected storage system is not registered for data collection.

Remote Replication

Managing remote replication sessions

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

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

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

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

The first option, bias, is a function of the two VMAX3 arrays taking part in the SRDF/Metro and is a required and integral component of the configuration. The second option, Witness, is an optional component of SRDF/Metro which allows a third VMAX or VMAX3 system to act as an external arbitrator to avoid an inconsistent result in cases where the bias functionality alone may not result in continued host availability of a surviving non-biased array.

SRDF

Configuration

Creating SRDF pairs

This procedure explains how to create SRDF pairs.

Before you begin:

- SRDF requires Enginuity version 5773-5876, or HYPERMAX OS 5977 or higher.
- 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 *Solutions Enabler Installation Guide*.
  - Create SRDF pairs containing standard and thin volumes, or thin and diskless volumes. To use this feature, the thin and diskless volumes must be on a storage system running Enginuity 5875 or higher, and the standard volume must be on a storage system running Enginuity 5875, 5773, or 5671.
- Meta volumes are supported on storage systems running Enginuity 5773 - 5876.
- On storage systems running HYPERMAX OS 5977 or higher, you can specify a RecoverPoint volume as the R1 volume.
- The cascaded R1 -> R21 -> R2 configuration of which an SRDF pair can be part, depends on the Enginuity/HYPERMAX OS version of each of the devices. The following combinations are supported:

R1	R21	R2
5977	5977	5977
5977	5977	5876
5977	5876	5876
5977	5876	5977
5876	5876	5977
5876	5977	5977
5876	5977	5876

To create an SRDF pair:

1. Select the storage 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 storage 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. The SRDF mode **Active Active** can only be selected if the SRDF Group is empty. If the SRDF Group already contains an Active Active session then the SRDF mode selector is not visible. For a list of possible modes, refer to SRDF group modes on page 515.

6. *Optional*: Select the **Adaptive Copy Mode**. Possible values are:
   - **Enabled: Disk Mode** — When set, the storage 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 storage system acknowledges all writes to the source (R1) volume as if it was a local volume. The new data accumulates in cache until it is successfully written to the source (R1) volume and the remote director has transferred the write to the target (R2) volume. This option requires that the source volume is on a storage system running Enginuity version 5773-5876.
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 storage system (**Local Start Volume**).
9. Type or select the number of the first volume in the range of volumes on the remote storage system (**Remote Start Volume**).
10. *Optional*: Depending on the **Mirror Type**, select to **Include Larger R2** volumes in the pair. or to **Include Smaller R1** volumes in the pair.
11. *Optional*: Click **Add to local Storage group**.
12. *Optional*: Select storage group to which the local SRDF devices will be added.
13. *Optional*: Click **Add to remote Storage group**.
14. *Optional*: Select storage group to which the remote devices will be added.
15. 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.
- When launching the **Create Pairs** dialog from a SRDF/Metro group, radio buttons indicating Witness or Bias will be displayed. If Witness is detected this radio button is defaulted to Witness. If Witness is not detected, the radio button group is disabled and defaulted to Bias.

2. Click **OK**.

Deleting SRDF pairs

Deleting SRDF pairs cancels the dynamic SRDF pairing by removing the pairing information from the storage system and converting the volumes from SRDF to regular volumes. This operation can be performed on a storage group, a SRDF/Metro, or a device group.

Half deleting SRDF pairs cancels the dynamic SRDF pairing information for one side (R1s or R2s) of the specified volume pairs and converts the volumes from RDF to regular volumes. This operation can only be performed on a device group.

If you select all pairs for a delete pair action, then the option to remove the devices from the device group, or the local or remote Storage Group is not displayed, as it will not render the device group, storage group, or SRDF/Metro unmanageable.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To delete SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups**, **Device Groups** or **SRDF/Metro**.
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 **Delete Pair** to open the **Delete Pair** dialog box.
   b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
   c. Select the **Half Delete** option if deleting one side of the volume pair.

d. *Optional*: Select **Remove from local Storage Groups**, **Remove from remote Storage Groups**, and **Remove from local Storage Groups** if the pair deletion results in devices that are no longer SRDF protected, and results in the related device groups becoming invalid.

e. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). 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 **Delete Pair** to open the **Delete Pair** dialog box.

c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

d. Select the **Half Delete** option if deleting one side of the volume pair.

e. *Optional*: Deselect the selected (by default) **Remove from Local Storage Groups**, **Remove from Remote Storage Groups**, and **Remove from Device Groups** check boxes. If you deselect the selected defaults, you will be warned if the pair deletion results in devices that are no longer SRDF protected, and results in the related device groups becoming invalid. This option is not displayed if all pairs are selected.

f. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Deleting SRDF pairs from details view

Moving SRDF pairs

This procedure explains how to move 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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To move SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
4. Select a group, click more >> , and select **Move** to open the **Move - Device Group** dialog box.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
6. Select **New SRDF Group**.
7. Select **Move Type**.
8. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Setting SRDF mode

This procedure explains how to set the mode of operation for an SRDF configuration. SRDF modes determine the following:

- How R1 volumes are remotely mirrored to R2 volumes across the SRDF links
- How I/OS are processed in an SRDF solution
- When acknowledgments are returned to the production host that issued a write I/O command

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

The **Adaptive Copy Mode** value **Enabled: WP Mode** is not available if the R1 mirror of an SRDF pair is on a storage system running HYPERMAX OS 5977, or higher.

To set SRDF mode:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 -**

Device Group dialog box.

b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

c. Select **SRDF Mode**, **Adaptive Copy Mode** and **AC Skew** to set the type of SRDF session modes (page 483).

d. Select **Use Consistent** to set consistent transition from asynchronous to synchronous mode.

e. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- **Show Advanced** to set the advanced SRDF session options (page 480). 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 483).

d. Select **Use Consistent** to set consistent transition from asynchronous to synchronous mode.

e. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Viewing SRDF volume pairs

This procedure explains how to view and manage the volume pairs in an SRDF group.

1. Select the storage system.
2. Select **Data Protection ›SRDF** to open the **SRDF** dashboard.
3. Select a device group from the list and click **View Details** to open the **SRDF Pair List** view.

   The following properties display:

   Show Group Details:

   Displays the following device group properties:

   - **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 storage system 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.

Local tab:

Displays the following local SRDF link properties:

- **Source Volume** — Source volume ID.
- **Source LDev** — Source logical volume ID
- **Group** — SRDF group ID.
- **Remote Symmetrix** — Remote storage system 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.
- **SRDF Mode** — SRDF copy type.
- **Local R1 Invalid** — Number of invalid R1 tracks on the source volume.
- **Local R2 Invalid** — Number of invalid R2 tracks on the source volume.
- **Remote R1 Invalid** — Number of invalid R1 tracks on the target volume.
- **Remote R2 Invalid** — Number of invalid R2 tracks on the target volume.

Hop2 tab:

Displays the following remote SRDF link properties:

- **Source LDev** — Source logical volume ID
- **Concurrent Volume** — Concurrent volume ID.
- **SRDF Group** — SRDF group ID.
- **Remote Symmetrix** — Remote storage system 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 461.
- **Split** — See Splitting SRDF pairs on page 477.
- **Suspend** — Suspending SRDF pairs on page 478
- **Restore** — See Restoring SRDF pairs on page 474.
- **View Details** — See Viewing SRDF volume pairs on page 447.
- **Resume** — Resuming SRDF links (page 470)
- **Failover** — Failing over (page 462)
- **Failback** — Failing back (page 463)
- **Activate SRDF/A** — See Controls on page 457.
- **Deactivate SRDF/A** — See Deactivating SRDF/A controls on page 458.
- **Invalidate** — See Invalidating R1/R2 volumes on page 465.
- **Ready** — See Making R1/R2 volumes ready on page 466.
- **Not Ready** — See Making R1/R2 volumes not ready on page 467.
- **R1 Update** — See Updating R1 volumes on page 479.
- **RW Enable** — See Read/write enabling R1/R2 volumes on page 469.
- **Write Disable** — See Write disabling R1/R2 volumes on page 470.
- **RW Disable R2** — See Read/write disabling R2 volumes on page 468.
- **Refresh** — See Refreshing R1 or R2 volumes on page 471.
- **Set Mode** — See Setting SRDF mode on page 446.

Viewing SRDF volume pair details

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Select a device group from the list and click **View Details** to open the **SRDF Pair List** view.
4. On the **Local** tab, select the pair and click **View Details** to open its details view.
5. Use the 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 storage system 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 461.
- **Split** — See Splitting SRDF pairs on page 477.
- **Suspend** — Suspending SRDF pairs on page 478
- **Restore** — See Restoring SRDF pairs on page 474.
- **Resume** — Resuming SRDF links (page 470)
- **Failover** — Failing over (page 462)
- **Failback** — Failing back (page 463)
- **Activate SRDF/A** — See Controls on page 457.
- **Deactivate SRDF/A** — See Deactivating SRDF/A controls on page 458.
- **Invalidate** — See Invalidating R1/R2 volumes on page 465.
- **Ready** — See Making R1/R2 volumes ready on page 466.
- **Not Ready** — See Making R1/R2 volumes not ready on page 467.
- **R1 Update** — See Updating R1 volumes on page 479.
- **RW Enable** — See Read/write enabling R1/R2 volumes on page 469.
- **Write Disable** — See Write disabling R1/R2 volumes on page 470.
- **RW Disable R2** — See Read/write disabling R2 volumes on page 468.
- **Refresh** — See Refreshing R1 or R2 volumes on page 471.
- **Set Mode** — See Setting SRDF mode on page 446.

Related Object 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.

Viewing SRDF protected storage group pairs

1. Select the storage system.
2. Select **Data Protection › SRDF**.
3. Click **Storage Groups** to open the **Storage Group** list view.
4. Select a storage group instance and click **View Details** to open the **Storage Group** pair list view.

   The following properties display, depending on the operating environment:

   - **Source Volume**— The name of the source volume.
   - **Source Type**— The source type of the source volume.
   - **SRDF Group**— The RDF group number.
   - **Remote Symmetrix**— The storage system serial number ID of the remote storage system.
   - **Target Volume**— The target volume ID.
   - **State** — The state of the storage group pair. Possible values are:
     - Consistent
     - Failed Over
     - Invalid
     - Partitioned
     - R1 Updated
     - R1 Update in progress
     - Suspended
     - Synchronization in progress
     - Synchronized
     - Transmit Idle

   If Unisphere detects an asynchronous state change event for a SRDF group from Solutions Enabler, it updates the Unisphere state for the SRDF group and its related SRDF device groups and SRDF storage groups. The **Storage Group** list view must be refreshed so that the latest state is reflected.

   - **Volume State** — The state of the source volume.
   - **Remote Volume State** — The state of the remote volume.
   - **SRDF Mode** — The SRDF copy mode.
   - **Local R1 Invalid**— The number of invalid R1 tracks on the source volume.

- **Local R2 Invalid**— The number of invalid R2 tracks on the source volume.
- **Remote R1 Invalid**— The number of invalid R1 tracks on the target volume.
- **Remote R2 Invalid** — The number of invalid R2 tracks on the target volume.

The following controls are available, depending on the operating environment:

- **Establish** — See Establishing SRDF pairs on page 461.
- **Split** — See Splitting SRDF pairs on page 477.
- **Suspend** — See Suspending SRDF pairs on page 478.
- **Restore** — See Restoring SRDF pairs on page 474.
- **View Details**— See Related Object panel on the facing page.
- **Resume** — See Resuming SRDF links on page 470.
- **Activate SRDF/A**— See Controls on page 457.
- **Deactivate SRDF/A**— See Deactivating SRDF/A controls on page 458.
- **Invalidate**— See Invalidating R1/R2 volumes on page 465.
- **Ready**— See Making R1/R2 volumes ready on page 466.
- **Not Ready**— See Making R1/R2 volumes not ready on page 467.
- **R1 Update**— See Updating R1 volumes on page 479.
- **RW Enable** — See Read/write enabling R1/R2 volumes on page 469.
- **Write Disable** — Write disabling R1/R2 volumes on page 470
- **RW Disable R2s**— Read/write disabling R2 volumes on page 468
- **Refresh** — Refreshing R1 or R2 volumes on page 471
- **Set Mode** — Setting SRDF mode on page 446

Viewing SRDF protected storage group pair properties

1. Select the storage system.
2. Select **Data Protection › SRDF**.
3. Click **Storage Groups** to open the **Storage Group** list view.
4. Select a storage group instance and click **View Details** to open the **Storage Group** pair list view.
5. Select a storage group pair instance and click **View Details** to open the **Storage Group** pair list properties view.

   The following properties display, depending on the operating environment:

   - **Storage Group**— The name of the storage group.
   - **Source Volume**— The name of the source volume.

- **SRDF Group**— The RDF group number.
- **Remote Symmetrix**— The storage system serial number ID of the remote storage system.
- **Remote SRDF Group**— The RDF group number.
- **Target Volume**— The target volume ID.
- **Pair State** — The state of the storage group pair.
- **Link Status**— Indicates link state.
- **SRDF Domino** — Indicates SRDF Domino state.
- **Source Volume SRDF State** — Indicates SRDF state of source volume.
- **Source Volume SRDF Type**—Indicates SRDF type of source volume.
- **Source Volume Track Size**—Indicates track size of source volume.
- **Target Volume SRDF State**— Indicates SRDF state of target volume.
- **Target Volume Track Size**— Indicates track size of target volume.

Related Object 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.

Viewing SRDF/Metro protected storage group pairs

Unisphere for VMAX provides SRDF monitoring and management for SRDF/Metro protected storage groups on storage systems running HYPERMAX OS 5977 or higher. Only single hop SRDF is supported for SRDF/Metro, that is, current or cascaded setups are not supported.

1. Select the storage system.
2. Select **Data Protection › SRDF**.
3. Click **SRDF/Metro** to open the SRDF/Metro storage group list view.
4. Select a storage group instance and click **View Details** to open SRDF/Metro storage group pair list view

   The following properties display, depending on the operating environment:

   - **Source Volume**— The name of the source volume.
   - **Source Type**— The source type of the source volume.
   - **SRDF Group**— The RDF group number.
   - **Remote Symmetrix**— The storage system serial number ID of the remote storage system.
   - **Target Volume**— The target volume ID.
   - **State** — The state of the storage group pair. Possible values are:

- ActiveActive
- ActiveBias
- Consistent
- Failed Over
- Invalid
- Partitioned
- R1 Updated
- R1 Update in progress
- Suspended
- Synchronization in progress
- Synchronized
- Transmit Idle

◆ **Volume State** — The state of the source volume.

◆ **Remote Volume State** — The state of the remote volume.

◆ **SRDF Mode** — The SRDF copy type.

The following controls are available, depending on the operating environment:

◆ **Delete Pair**— See Deleting SRDF pairs on page 459.

◆ **View Details**— See "Related Object panel" on page 457

Viewing SRDF Metro protected storage group pair properties

Unisphere for VMAX provides SRDF monitoring and management for SRDF/Metro protected storage groups on storage systems running HYPERMAX OS 5977 or higher. Only single hop SRDF is supported for SRDF/Metro, that is, current or cascaded setups are not supported.

1. Select the storage system.
2. Select **Data Protection › SRDF**.
3. Click **SRDF/Metro** to open the SRDF/Metro storage group list view.
4. Select a storage group instance and click **View Details** to open SRDF/Metro storage group pair list view
5. Select a storage group pair instance and click **View Details** to open the SRDF/Metro storage group pair list properties view.

   The following properties display, depending on the operating environment:

   ◆ **Storage Group**— The name of the storage group..

   ◆ **Source Volume**— The name of the source volume.

   ◆ **SRDF Group**— The RDF group number.

   ◆ **Remote Symmetrix**— The storage system serial number ID of the remote storage system.

   ◆ **Remote SRDF Group**— The RDF group number.

- **Target Volume**— The target volume ID.
- **Pair State** — The state of the storage group pair.
- **SRDF Mode** — The SRDF copy type.

Related Object 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.

Controls

Activating SRDF/A controls

This procedure explains how to activate SRDF/A control actions (page 513) 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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To activate SRDF/A controls:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups).
   c. Select **Activate Type**.
   d. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Deactivating SRDF/A controls

This procedure explains how to deactivate SRDF/A control actions (page 513) that detect cache overflow conditions and take corrective action.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To deactivate SRDF/A controls:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups).

   c. Select **Deactivate Type**.

   d. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to set the advanced options as described in SRDF session options on page 480. 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. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced options as described in SRDF session options on page 480. Select the advanced options and click **OK**.

.

Deleting SRDF pairs

Deleting SRDF pairs cancels the dynamic SRDF pairing by removing the pairing information from the storage system and converting the volumes from SRDF to regular volumes. This operation can be performed on a storage group, a SRDF/Metro, or a device group.

Half deleting SRDF pairs cancels the dynamic SRDF pairing information for one side (R1s or R2s) of the specified volume pairs and converts the volumes from RDF to regular volumes. This operation can only be performed on a device group.

If you select all pairs for a delete pair action, then the option to remove the devices from the device group, or the local or remote Storage Group is not displayed, as it will not render the device group, storage group, or SRDF/Metro unmanageable.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To delete SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups**, **Device Groups** or **SRDF/Metro**.
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 **Delete Pair** to open the **Delete Pair** dialog box.

   b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

   c. Select the **Half Delete** option if deleting one side of the volume pair.

   d. *Optional*: Select **Remove from local Storage Groups**, **Remove from remote Storage Groups**, and **Remove from local Storage Groups** if the pair deletion results in devices that are no longer SRDF protected, and results in the related device groups becoming invalid.

e. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). 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 **Delete Pair** to open the **Delete Pair** dialog box.

c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

d. Select the **Half Delete** option if deleting one side of the volume pair.

e. *Optional*: Deselect the selected (by default) **Remove from Local Storage Groups**, **Remove from Remote Storage Groups**, and **Remove from Device Groups** check boxes. If you deselect the selected defaults, you will be warned if the pair deletion results in devices that are no longer SRDF protected, and results in the related device groups becoming invalid. This option is not displayed if all pairs are selected.

f. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Deleting SRDF pairs from the SRDF List Volumes View

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To delete SRDF pairs from the SRDF List Volumes View

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools** to open the **SRDF** dashboard.
3. Click **SRDF Groups**.
4. Select a group and click **View Details**.
5. Click the Volumes related object.
6. Select a volume and click and select **Delete Pairs** to open the **Delete Pair** dialog box.
7. Select the **Half Delete** option if deleting one side of the volume pair.

8. *Optional*: Deselect the selected (by default) **Remove from Local Storage Groups**, **Remove from Remote Storage Groups**, and **Remove from Device Groups**check boxes. If you deselect the selected defaults, you will be warned if the pair deletion results in devices that are no longer SRDF protected, and results in the related device groups becoming invalid. This option is not displayed if all pairs are selected.
9. *Optional*: Select **Use Force**.
10. Click **OK**.

Establishing SRDF pairs

This procedure explains how to start SRDF pair mirroring.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

You can run an establish operation on a cascaded R1 -> R21 -> R2 configuration if any of the storage systems in the cascaded configuration is running HYPERMAX OS Q1 2015 SR or later.

To establish SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups**, **Device Groups** or **SRDF/Metro**.
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 **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 (only applicable for device groups) .
   d. Select **Witness** or **Bias** (only applicable for SRDF/Metro). Witness, if available, is the default option. If Witness is not available, Bias is set by the system and the radio buttons are disabled.
   e. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced options as described in SRDF session options on page 480. 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. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced options as described in SRDF session options on page 480. Select the advanced options and click **OK**.

Failing over

In a period of scheduled downtime for maintenance, or after a serious system problem which has rendered either the host or storage 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.

Before you begin:

If the target (R2) volume is on a storage system running HYPERMAX OS 5977 or higher, and the mode of the source (R1) volume is Adaptive Copy Write Pending, SRDF will set the mode to Adaptive Copy Disk.

As a result of a failover (with establish or restore) operation, a cascaded R1 -> R21 -> R2 configuration can be created if any of the storage systems in the cascaded configuration is running HYPERMAX OS Q1 2015 SR or later.

To initiate a failover:

1. Select the storage system.
2. Select **Data Protection > SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **Failover** to open the **Failover** dialog box.

   b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

   c. Select the fail over SRDF session options (page 480).

   d. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to set the advanced SRDF session options (page 480).

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 480).

d. Do one of the following::

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To initiate a failback:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **Failback** to open the **Failback** dialog box.

   b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

   c. Select any of the available session options, as described in SRDF session options on page 480.

   d. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

- Click **Show Advanced** to set the advanced SRDF session options (page 480). 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 480.

d. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Half-swapping SRDF/Metro pairs

When the SRDF device pairs of an SRDF/Metro configuration are Not Ready (NR) on the link, and the SRDF pair state is Partitioned, a half swap operation is allowed. If the half swap is issued to the R2, the SRDF link to the R1 must be unavailable.

If the half swap is issued to the R1, the SRDF link to the other side must be available and the SRDF pair must be seen as R1 – R1 (duplicate pair).

Before you begin:

SRDF requires HYPERMAX OS 5977 or higher.

To half swap SRDF/Metro pair:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **SRDF/Metro**.
4. Select a group, click more >> , and select **Half Swap** to open the **Half Swap** dialog box.
5. Click one of the following:
   - **OK** to start the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To invalidate R1/R2 volumes:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups.
   d. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). Select the

advanced options and click **OK**.

Making R1/R2 volumes ready

This procedure explains how to set the R1 (source) or R2 (target) volumes ready to their local hosts.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To make R1/R2 volumes ready:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups).
   d. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To make R1/R2 volumes not ready:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups).
   d. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced SRDF session options (page 480).

Select the advanced options and click **OK**.

Read/write disabling R2 volumes

This procedure explains how to read/write disable target (R2) volumes to their local hosts.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To read/write disable R2 volumes:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups).
   c. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To read/write enable R1/R2 volumes:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 (only applicable for device groups).
   c. Select **RW Enable R1s** or **RW Enable R2s** volume type.
   d. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To resume SRDF links:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **Resume** to open the **Resume** dialog box.
   b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
   c. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Write disabling R1/R2 volumes

This procedure explains how to write disable source (R1) volumes/target (R2) volumes to their local hosts.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To write disable R1/R2 volumes:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.

   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** to open the dialog box.

   b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

   c. Select **Write Disable R1s** or **Write Disable R2s** volume type.

   d. Do one of the following:

   - Do **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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:

   - Do **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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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:

- SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.
- 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 storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
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 **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 480). 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 480). 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**.

Resetting original device identity

After deleting a SRDF/Metro pair, the unbiased devices keep the new identity. This procedure explains how to reset the original device identity.

Before you begin:

SRDF requires HYPERMAX OS 5977 or higher.

To establish SRDF pairs:

1. Select the storage system.
2. Do the following, depending on whether you want to perform the operation at the group level or pair level:
   Group level:
   a. From the Data Protection dashboard:
      - Select **Data Protection › Protection Dashboard** to open the **Protection** dashboard.
      - Click **Total** to view the list of all storage groups on the storage system.
      - Select a former unbiased SRDF/Metro storage group and click **Reset SRDF/Metro Identity** to open the **Reset Original Identity** dialog box.

   b. From the Storage dashboard:
      - Select **Storage › Storage Groups Dashboard** to open the **Storage Groups** dashboard.
      - Click **Total** to view the list of all storage groups on the storage system.
      - Select a former unbiased SRDF/Metro storage group and click **Reset SRDF/Metro Identity** to open the **Reset Original Identity** dialog box.

   c. Do one of the following:
      - 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
      - Click **Show Advanced** to set the advanced options as described in SRDF session options on page 480. Select the advanced options and click **OK**.

   Pair level:
   a. Select **Storage › Volumes** to open the **Storage Volumes** view.
   b. Filter the view to display volume(s) that were formally part of a SRDF/Metro pair.
   c. Do one of the following:
      - Select a volume and click **Reset SRDF/Metro Identity** to open the **Reset Original Identity** dialog box.

- Select a volume, click **View Details** to view the volume details, and then click **Reset SRDF/Metro Identity** to open the **Reset Original Identity** dialog box.

d. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Show Advanced** to set the advanced options as described in SRDF session options on page 480. Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To restore SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups**, **Device Groups** or **SRDF/Metro**.
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 **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 (only applicable for device groups).
   c. Select **Full** or **Incremental** restore type.
   d. Select **Witness** or **Bias** (only applicable for SRDF/Metro). Witness, if available, is the default option. If Witness is not available, Bias is set by the system and the radio buttons are disabled.
   e. Do one of the following:
      - **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

- **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Setting bias location

This procedure explains how to set bias. If bias is chosen to be set as part of the **Suspend** operation, the side with the bias is the side that the host can see after the **Suspend** operation completes.

To set bias:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Select **SRDF/Metro**.
4. Select a group, click more >> , and select **Set Bias Location** to open the dialog box.
5. View the **Move Bias** setting that moves the bias from one side to the other (only applicable for SRDF/Metro). This will take effect after this operation is executed.The bias indicates the side that remains visible to the host.
6. Do one of the following:
   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - Click **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Setting the SRDF GCM flag

This procedure supports the setting of the SRDF GCM flag both from storage Group perspective and at individual device level.

Before you begin:

SRDF requires HYPERMAX OS 5977 or higher.

To establish SRDF pairs:

1. Select the storage system.
2. Do the following, depending on whether you want to perform the operation at the group level or pair level:
   Group level:
   a. From the Data Protection dashboard:
      - Select **Data Protection › Protection Dashboard** to open the **Protection**

dashboard.

- Click **Total** to view the list of all storage groups on the storage system.
- Select a storage group and click **Set SRDF GCM** to open the **Set GCM** dialog box.
- Click **OK**.

b. From the Storage Volumes View:

- Select **Storage › Storage Volumes** to open the **Storage Volumes** view.
- Select a storage group and click **Set SRDF GCM** to open the **Set GCM** dialog box.
- Click **OK**.

c. From the Storage Volumes Details View:

- Select **Storage › Storage Volumes** to open the **Storage Volumes** view.
- Select a storage group and click **View Details** to open the**Storage Volumes** view.
- Click **Set SRDF GCM** to open the **Set GCM** dialog box.
- Click **OK**.

Setting volume status

After deleting an SRDF/Metro pair, the volumes can be in a Not Ready state. This dialog allows you to set the volume state.

Before you begin:

SRDF requires HYPERMAX OS 5977 or higher.

To set the volume state:

1. Select the storage system.
2. Do the following, depending on whether you want to perform the operation at the group level or pair level:
   Group level:

   a. From the Data Protection dashboard:

   - Select **Data Protection › Protection Dashboard** to open the **Protection** dashboard.
   - Click **Total** to view the list of all storage groups on the storage system.
   - Select a former unbiased SRDF/Metro storage group and click **Set Volume Status** to open the **Set Volume Status** dialog box.

   b. From the Storage dashboard:

   - Select **Storage › Storage Groups Dashboard** to open the **Storage Groups** dashboard.
   - Click **Total** to view the list of all storage groups on the storage system.

- Select a former unbiased SRDF/Metro storage group and click **Set Volume Status** to open the **Set Volume Status** dialog box.

c. Click **OK**.

Splitting SRDF pairs

This procedure explains how to stop SRDF pair mirroring.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To split SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.

   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 (only applicable for device groups).

   c. Select **Use Immediate** for immediate split on asynchronous devices.

   d. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

Suspending SRDF pairs

This procedure explains how to stop data transfer between SRDF pairs.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To suspend SRDF pairs:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups**, **Device Groups** or **SRDF/Metro**.

   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** dialog box.

   b. Select the **Hop2** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).

   c. Select **Use Immediate**or **Use Consistency Exempt** SRDF session options (page 480).

   d. Click **Move Bias** to move the bias from one side to the other (only applicable for SRDF/Metro). The side with the bias set is the side that the host can see after the suspend action completes.

   e. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

   Pair level:

   a. Select a storage or device 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. Do one of the following:

   - 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 851 and Previewing on page 850.
   - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To swap SRDF personalities:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.
4. Select a group, click more >> , and select **Swap** to open the **Swap** dialog box.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
6. For optional Refreshing R1 or R2 volumes (page 471), select **R1, R2 or None**.
7. For optional half swapping, select **Half Swap**.
8. Click one of the following:
   - **OK** to start the operation now.
   - **Show Advanced** to set the advanced SRDF session options (page 480). 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.

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To update R1 volumes:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Storage Groups** or **Device Groups**.

   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 (only applicable for device groups).

c. Select **Remote** if R1 volumes are a remote.

d. Click one of the following:

- **Show Advanced** to set the advanced SRDF session options (page 480). 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. Do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- **Show Advanced** to set the advanced SRDF session options (page 480). Select the advanced options and click **OK**.

SRDF session options

Session option	Description	Available with action
Bypass	Bypasses the exclusive locks for the local and/or remote storage system during SRDF operations. Use this option only if you are sure that no other SRDF operation is in progress on the local and/or remote storage systems.	Establish Failback Failover Restore Incremental Restore Split Suspend Swap Write Disable R1 Ready R1 Ready R2

Session option	Description	Available with action
		RWDisableR2 Enable Disable
Consistent	Allows only consistent transition from async to sync mode.	Activate
Consistency Exempt	Allows you to add or remove volumes from an RDF group that is in Async mode without requiring other volumes in the group to be suspended.	Half Move Move Suspend
Establish	Fails over the volume pairs, performs a dynamic swap, and incrementally establishes the pairs. This option is not supported when volumes operating in Asynchronous mode are read/write on the RDF link. To perform a fail over operation on such volumes, specify the Restore option detailed higher in this table.	Failover
Force	Overrides any restrictions and forces the operation, even though one or more paired volumes may not be in the expected state. Use caution when checking this option because improper use may result in data loss.	Establish Incremental Establish Restore Incremental Restore Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable Disable Swap
Immediate	Causes the suspend, split, and failover actions on asynchronous volumes to happen immediately.	Suspend Split Failover
NoWD	No write disable - bypasses the check to ensure that the target of operation is write disabled to the host. This applies to the source (R1) volumes when used with the Invalidate R1option and to the target (R2) volumes when used with the Invalidate R2 option.	
SymForce	Forces an operation on the volume pair including pairs that	Restore

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

Session option	Description	Available with action
	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). SRDFR/Star is supported on Enginuity 5671 - 5876. The *Solutions Enabler SRDF Family CLI Product Guide* contains more information on SRDF/Star.	Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable 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. This mode requires that the device group containing the RDF pairs with R1 mirrors be on a storage system running Enginuity 5773-5876.
Synchronous	Provides the host access to the source (R1) volume on a write operation only after the storage system containing the target (R2) volume acknowledges that it has received and checked the data.
Asynchronous	The storage system acknowledges all writes to the source (R1) volumes as if they were local 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. For storage systems running Enginuity 5875, you can put an RDF relationship into Asynchronous mode when the R2 device is a snap source volume.
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.

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.

Note Virtual LUN migration requires Enginuity 5874 - 5876.

Virtual LUN Migration performs tiered storage migration by moving data from one RAID group to another, or from one thin pool to another. It is also fully interoperable with all other 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 Enginuity 5874 - 5876.
- 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 storage 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 long and are case sensitive.
5. Select a **Migration Type**. Available options are **Disk Group** and **Virtual Pool**.
6. Select a **Target Disk Group**.
7. Select the RAID **Protection type**.
8. Select **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** to create the migration session.

Migrating regular volumes

This procedure explains how to migrate individual regular volumes.

Before you begin:

- Virtual LUN migration requires Enginuity 5874 - 5876.
- To migrate volumes when access control is enabled, you must grant Symmetrix Optimizer access rights to the volumes.

To migrate regular volumes:

1. Select the storage 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 wether 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 - 5876.
- 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 storage 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. *Optional:* Expand **Show Advanced**. From the **Migrate allocations from pool** menu, select a pool from which to migrate allocations.
8. Click **OK**.

Migrating thin volumes

This procedure explains how to migrate individual thin volumes.

Before you begin:

- Virtual LUN migration requires Enginuity 5874 - 5876.
- 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 storage 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 wether to **Pin volumes** so that they cannot be moved by any FAST automated process.
8. *Optional:* Expand **Show Advanced**. From the **Migrate allocations from pool** menu, select a pool from which to migrate allocations.
9. Click **OK**.

Viewing VLUN migration sessions

1. Select the storage system.
2. Select **Data Protection › Virtual LUN Migration** to open the **Virtual LUN 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 280
   - **View Details** — Viewing VLUN migration session details below

Viewing VLUN migration session details

1. Select the storage 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

Before you begin:

SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To create an SRDF/A DSE pool:

1. Select the storage 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

Before you begin:

- SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To delete an SRDF/A DSE pool:

1. Select the storage 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

Before you begin:

- SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To add a volume to an SRDF/A DSE pool:

1. Select the storage 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

Before you begin:

- SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To enable all volumes in an SRDF/A DSE pool:

1. Select the storage 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

Before you begin:

- SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To disable all volumes in an SRDF/A DSE pool:

1. Select the storage 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

Before you begin:

- SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To view an SRDF/A DSE pool:

1. Select the storage 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 DSE Pool** —See Replication Groups and Pools on page 489.
   - **Add** — See Adding volumes to SRDF/A DSE pools on the previous page.
   - **Delete** — See Deleting SRDF/A DSE pools on page 489.
   - **View Details** — See Viewing SRDF DSE pool details below.
   - **Enable All** — See Enabling all volumes in SRDF/A DSE pools on the previous page.
   - **Disable All** — See Disabling all volumes in SRDF/A DSE pools on the previous page.
   - **Assign Dynamic Cache Partition** — Assigning dynamic cache partitions on page 875

Viewing SRDF DSE pool details

Before you begin:

- SRDF/A DSE pools are supported on storage systems running Enginuity 5773 - 5876.

To view the details of an SRDF/A DSE pool:

1. Select the storage 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** — Storage 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 (GB)** — Sum of all enabled and disabled volumes in the pool.
- **Enabled Capacity (GB)** — Sum of all enabled volumes in the pool.
- **Free Capacity (GB)** — 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 DSE Pool** — See Replication Groups and Pools on page 489.
- **Add** — See Adding volumes to SRDF/A DSE pools on page 490.
- **Delete DSE Pool** — See Deleting SRDF/A DSE pools on page 489.
- **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 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 Snap Pool** — See TimeFinder Snap Pools on page 493.
- **Add** — See Adding volumes to TimeFinder/Snap pools on page 493..
- **View Details** — See Viewing TimeFinder/Snap pool details below.
- **Delete** — See Deleting TimeFinder/Snap Pools on the previous page.
- **Enable All** — See Enabling all volumes in TimeFinder/Snap pools on the previous page.
- **Disable All** — See Disabling all volumes in TimeFinder/Snap pools on the previous page.
- **Assign Dynamic Cache Partition** — Assigning dynamic cache partitions on page 875

Viewing TimeFinder/Snap pool details

1. Select the storage 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 (GB)** — Sum of all enabled and disabled volumes in the pool.
- **Enabled Capacity (GB)** — Sum of all enabled volumes in the pool.
- **Free Capacity (GB)** — 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 Snap Pool** — See TimeFinder Snap Pools on page 493..
- **Add** — See Adding volumes to TimeFinder/Snap pools on page 493.
- **Delete** — See Deleting TimeFinder/Snap Pools on page 494.
- **Enable All** — See Enabling all volumes in TimeFinder/Snap pools on page 494.
- **Disable All** — See Disabling all volumes in TimeFinder/Snap pools on page 494.
- **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 storage 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 storage systems within the fabric topology.

Before you begin:

- SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.
- The maximum number of supported RDF groups differs by Enginuity version:

Enginuity	Maximum number of RDF Groups supported			Group numbers
	per storage system	per director	per port	
5977 or higher	250	250	250	1 to 250
5874 or higher	250	64	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 storage system and number of RDF groups defined on a director is determined solely by the Enginuity level running on the storage system where the group is being defined.
- When specifying a local or remote director for a VMAX system running HYPERMAX OS 5977, you can select one or more SRDF ports.

To create an SRDF group:

1. Select the storage system.
2. Select **Data Protection › Common Tasks › Create SRDF groups** to open the **Create SRDF Group** dialog box.
3. Select a **Communication Protocol** to use when moving data across the SRDF links. The value you select here will populate the **Director** list.
4. Type an **SRDF Group Label** (name).
5. Type an **SRDF Group Number**.
6. Select the local **Director** through which the group will communicate.
7. *Optional*: To refresh the remote storage system information, click **Scan**.
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 **SRDF/Metro Witness Group**. This checkbox is selectable when the local Symmetrix and the remote selected Symmetrix are both Witness capable.
12. 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 storage 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 storage 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

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To modify SRDF groups:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › SRDF groups** to open the **SRDF groups** list view.
3. Select a group and click **Edit** to open the **Edit SRDF group** dialog box.
4. Do any number of the following steps:
   a. Select a new local **Director** through which the group will communicate.

      When specifying a local or remote director for a VMAX system running HYPERMAX OS 5977, you can select one or more SRDF ports.

   b. Select a new **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 storage system.
2. Select **Data Protection › Replication Groups and 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 1 to 60 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** 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 storage system.
2. Select **Data Protection › Replication Groups and 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.
   For systems running HYPERMAX OS 5977, this option may not be available.
5. Type the percentage of the storage system's write pending limit (**Threshold**). Once the cache usage of all active groups in the storage 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.
   For systems running HYPERMAX OS 5977, **Autostart** is always enabled.

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 Enginuity 5773.150, this feature is supported with thin devices.
8. Click **OK**.

Setting SRDF/A pace attributes

1. Select a storage system.
2. Select **Data Protection › Replication Groups and 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 Pacing 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 volumes, the source (R1) volumes become target (R2) volumes and the target (R2) volumes become source (R1) volumes.

Before you begin:

- If the target (R2) volume is on a storage system running HYPERMAX OS 5977 or higher, and the mode of the source (R1) volume is Adaptive Copy Write Pending, SRDF will set the mode to Adaptive Copy Disk.
- As a result of a swap, or half swap, operation, a cascaded R1 -> R21 -> R2 configuration can be created if any of the storage systems in the cascaded configuration is running HYPERMAX OS Q1 2015 SR or later.

To swap SRDF groups:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and 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.
5. *Optional*: Indicate if a half swap is required.
6. *Optional*: Select to Start refreshing the selected mirror (**Start Copy**).
7. If the RDF group you are swapping is part of an SRDF/Star configuration, select **Star**.
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 851 and Previewing on page 850.joepedusjj
   - 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 storage 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 storage 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 480). Select the advanced options and click **OK**.

Deleting SRDF groups

Before you begin:

SRDF requires Enginuity version 5773 - 5876, or HYPERMAX OS 5977 or higher.

To delete SRDF groups:

1. Select the storage system.
2. Select **Data Protection › Replication Groups & Pools › SRDF groups** to open the **SRDF Groups** list view.
3. Select the SRDF group, click more >> , and select **Delete**.
4. Click **OK** in the confirmation message.

Viewing SRDF groups

1. Select the storage system.
2. Select **Data Protection › Replication Groups and 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** — SRDF group label, for example, Async, Metro, Witness.
- **Remote SRDF group** — Remote RDF group number.
- **Remote Symmetrix** — Remote Symmetrix serial ID. This field only displays when the SRDF pair state is Partitioned.
- **Type** — Type of group, for example, Dynamic or Witness.
- **SRDF Modes** — SRDF modes associated with the SRDF group.
- **SRDF group Flags** — SRDF group flags. This field only applies/appears for Symmetrix systems running Enginuity 5773-5876.
- **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 RDFA flags on page 514.
- **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 496.
- **Create Pairs** — See SRDF on page 441.
- **Edit** — See Modifying SRDF groups on page 498.
- **View Details** — See Viewing SRDF group details on the facing page.
- **Assign Dynamic Cache Partition** — See Assigning dynamic cache partitions on page 875.
- **Delete** — See Deleting SRDF groups on the previous page.
- **Swap Groups** — See Swapping SRDF groups on page 501.
- **SRDF/A DSE Settings** — See Setting SRDF/A DSE attributes on page 500.
- **SRDF/A Pacing Settings** — See Setting SRDF/A pace attributes on page 501.
- **SRDF/A Settings** — See Setting SRDF/A group attributes on page 500.

Viewing SRDF group details

1. Select the storage 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 **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 properties listed depend on the specifics of the storage system. Some or all of the following properties display:

- **SRDF Group Number** — SRDF group number.
- **SRDF Group Label** — SRDF group label.
- **SRDF director number(s)** — SRDF director numbers.
- **Director Identity** — Director identifier(s).
- **Remote SRDF Group number(s)** — Remote group number(s).
- **Remote Symmetrix** — Remote storage system serial ID(s).
- **Remote SRDF director number(s)** — Remote SRDF director numbers.
- **Remote Director Identity** — Remote director identifier(s).
- **SRDF Modes**— SRDF Modes. Possible values are: N/A, Adaptive Copy, Synchronous, Asynchronous, Active.
- **Online**—Online. Possible values are: Yes, No.
- **SRDF/Metro**SRDF/Metro. Possible values are: Yes, No.
- **SRDF/Metro Witness Degraded**—SRDF/Metro Witness Degraded. Possible values are: Yes, No.
- **SRDF Group Flags** — SRDF group flags. This field only applies/appears for storage systems running Enginuity 5773-5876.
- **Prevent Auto Link Recovery** — Indicates the state of preventing automatic data copy across SRDF links upon recovery.
- **Copy Jobs** — Maximum number of SRDF copy jobs per SRDF group.
- **Prevent RAs Online Upon Power On** — Indicates the state of preventing the SRDF. 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.

- **SRDF Group Configuration** — RA group configuration. Possible values are: Dynamic, Static, Witness.
- **Link Limbo (sec)** — Number of seconds (0-10) for the storage system to continue checking the local SRDF link status.
- **Minimum Cycle Time** — Minimum cycle time (seconds) configured for this session.
- **Duration of last transmit cycle** — Duration of the last transmit cycle.
- **Transmit Queue Depth on R1 side** —Depth of the transmit queue on the source side.
- **Transmit Idle Time** — Time the transmit cycle has been idle.
- **Transmit Idle** — Whether SRDF/A Transmit Idle state is active for the SRDF 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 Mode** — The SRDF/A mode. The status of the property can be Single-session, MSC, or N/A.
- **MSC Cleanup Required** —Indicates if MSC cleanup is required. The status of the property can be Yes, No, or N/A.
- **SRDF/A Session Status** —The SRDF/A session status. The status of the property can be Active, Inactive, or N/A.
- **SRDF/A Consistency Protection** —Indicates if consistency protection is enabled. The status of the property can be Enabled, Disabled, or N/A.
- **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 storage systems write pending limit.
- **SRDF/A Write Pacing Status** — Indicates if SRDF/A write pacing is active.
- **SRDF/A Contains Pace Capable Volumes** — 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.
- **Device Pacing Supported** — Indicates if SRDF/A device pacing is supported.
- **Group Level Pacing State** — Indicates if group level write pacing is enabled or disabled.
- **Device Pacing Activated** — Group-level pacing status of the SRDF/A session. The status of the feature can be Active, Inactive, N/A.
- **Group Pacing Auto Start** — Indicates if group pacing auto start is enabled/disabled on the SRDF group.
- **Device Pacing Auto Start** — Indicates if device pacing auto start is enabled/disabled on the SRDF group.
- **SRDF Software Compression** — Indicates if software compression is enabled/disabled on the SRDF group.
- **RDF Single Round Trip** — Indicates if single round trip is enabled/disabled on the SRDF group.
- **RDF Hardware Compression** — Indicates if hardware compression is enabled/disabled on the SRDF group.
- **SRDF Software Compression Support** — Indicates if SRDF software compression is enabled or disabled.
- **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 storage system.
- **SRDF Group Hardware Compression** — Indicates if SRDF hardware compression is enabled or disabled for the SRDF group.
- **Star Mode** — Indicates if SRDF group is in a star configuration.
- **SQAR Mode** — Indicates if SRDF group is in a SQAR configuration.

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.

Viewing SRDF protected device groups

The **SRDF** dashboard provides you with a single place to monitor and manage SRDF sessions on a storage system. This includes device groups types R1, R2, and R21. See "Managing remote replication sessions" on page 441for additional information.

Before you begin:

- SRDF requires Enginuity version 5773-5876, or HYPERMAX OS 5977.
- The following configurations are not supported:
  - An R21 or R22 SRDF device on a system running HYPERMAX OS 5977.
  - A cascaded SRDF configuration containing a system running HYPERMAX OS 5977.
  - A concurrent R22 configuration containing a system running HYPERMAX OS 5977.

To access the SRDF dashboard:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Device Groups** to open the Device Group list view.

SRDF dashboard:

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 the group is valid or invalid for SRDF management.

The following controls are available:

- **Establish** — See Establishing SRDF pairs on page 461.
- **Split** — See Splitting SRDF pairs on page 477.
- **Suspend** — See Suspending SRDF pairs on page 478.
- **Restore** — See Restoring SRDF pairs on page 474.
- **View Storage Groups** — Viewing storage groups (Enginuity 5874 - 5876) on page 155
- **View Details** — See Viewing SRDF volume pairs on page 447.
- **Resume** — Resuming SRDF links on page 470
- **Failover** — Failing over on page 462
- **Failback** — Failing back on page 463
- **Swap** — See Swapping SRDF personalities on page 479.
- **Enable Consistency** — Enabling consistency protection on page 502

- **Disable Consistency** — See Disabling consistency protection on page 503.
- **Delete Pair** — Deleting SRDF pairs on page 459
- **Move** — Deleting SRDF pairs from details view on page 445
- **Activate SRDF/A** — See Controls on page 457.
- **Deactivate SRDF/A** — See Deactivating SRDF/A controls on page 458.
- **Invalidate** — See Invalidating R1/R2 volumes on page 465.
- **Ready** — See Making R1/R2 volumes ready on page 466.
- **Not Ready** — See Making R1/R2 volumes not ready on page 467.
- **R1 Update** — See Updating R1 volumes on page 479.
- **RW Enable** — See Read/write enabling R1/R2 volumes on page 469.
- **Write Disable** — See Write disabling R1/R2 volumes on page 470.
- **RW Disable R2s** — See Read/write disabling R2 volumes on page 468.
- **Refresh** — See Refreshing R1 or R2 volumes on page 471.
- **Set Mode** — See Setting SRDF mode on page 446.

Viewing SRDF Metro protected storage groups

Unisphere for VMAX 8.1 and higher provides SRDF monitoring and management for SRDF Metro protected storage groups. Only single hop SRDF is supported for SRDF/metro, that is, current or cascaded setups are not supported. See See "Managing remote replication sessions" on page 441for additional information.

1. Select the storage system.
2. Select **Data Protection › SRDF**.
3. Click **SRDF/Metro** to open the SRDF/Metro storage group list view.

   The following properties display, depending on the operating environment:

   - **Storage Group**— User-defined storage group name.
   - **Capacity(GB)** — Total capacity of the storage group in GB.
   - **State** — The state of the storage group. Possible values are:
     - Active Active
     - Consistent
     - Failed Over
     - Invalid
     - Partitioned
     - R1 Updated
     - R1 Update in progress
     - Suspended
     - Synchronization in progress
     - Synchronized
     - Transmit Idle

   If Unisphere detects an asynchronous state change event for a SRDF group from Solutions Enabler, it updates the Unisphere state for the SRDF group and its related SRDF device groups and SRDF storage groups. The **Storage Group** list view must be refreshed so that the latest state is reflected.

   - **SRDF Mode** — The SRDF copy type.
   - **Local Bias**—Bias refers to the side (R1 or R2) of the SRDF group that should remain ready (RW) to the host(s) if the link becomes Not Ready (NR). The non-bias side is made NR to the host(s) in this case.
   - **SRDF Group**— The RDF group number.
   - **Masking Views**— The name of the masking group.
   - **Emulation** — The emulation type (ALL, FBA, CKD).

   The following controls are available, depending on the operating environment:

   - **Establish** — See Establishing SRDF pairs on page 461.
   - **Suspend** — See Suspending SRDF pairs on page 478.
   - **Restore** — See Restoring SRDF pairs on page 474.

- **View Storage Groups**— See Viewing storage group details (HYPERMAX OS 5977 or higher) on page 150.
- **View Details**— See Viewing SRDF/Metro protected storage group pairs on page 455.
- **Set Bias**— See "Setting bias location" on page 475
- **Delete Pair**— See Deleting SRDF pairs on page 459.

Viewing SRDF protected device groups

The **SRDF** dashboard provides you with a single place to monitor and manage SRDF sessions on a storage system. This includes device groups types R1, R2, and R21. See See "Managing remote replication sessions" on page 441for additional information.

Before you begin:

- SRDF requires Enginuity version 5773-5876, or HYPERMAX OS 5977.
- The following configurations are not supported:
  - An R21 or R22 SRDF device on a system running HYPERMAX OS 5977.
  - A cascaded SRDF configuration containing a system running HYPERMAX OS 5977.
  - A concurrent R22 configuration containing a system running HYPERMAX OS 5977.

To access the SRDF dashboard:

1. Select the storage system.
2. Select **Data Protection › SRDF** to open the **SRDF** dashboard.
3. Click **Device Groups** to open the Device Group list view.

SRDF dashboard:

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 the group is valid or invalid for SRDF management.

The following controls are available:

- **Establish** — See Establishing SRDF pairs on page 461.
- **Split** — See Splitting SRDF pairs on page 477.
- **Suspend** — See Suspending SRDF pairs on page 478.

- **Restore** — See Restoring SRDF pairs on page 474.
- **View Storage Groups** — Viewing storage groups (Enginuity 5874 - 5876) on page 155
- **View Details** — See Viewing SRDF volume pairs on page 447.
- **Resume** — Resuming SRDF links on page 470
- **Failover** — Failing over on page 462
- **Failback** — Failing back on page 463
- **Swap** — See Swapping SRDF personalities on page 479.
- **Enable Consistency** — Enabling consistency protection on page 502
- **Disable Consistency** — See Disabling consistency protection on page 503.
- **Delete Pair** — Deleting SRDF pairs on page 459
- **Move** — Deleting SRDF pairs from details view on page 445
- **Activate SRDF/A** — See Controls on page 457.
- **Deactivate SRDF/A** — See Deactivating SRDF/A controls on page 458.
- **Invalidate** — See Invalidating R1/R2 volumes on page 465.
- **Ready** — See Making R1/R2 volumes ready on page 466.
- **Not Ready** — See Making R1/R2 volumes not ready on page 467.
- **R1 Update** — See Updating R1 volumes on page 479.
- **RW Enable** — See Read/write enabling R1/R2 volumes on page 469.
- **Write Disable** — See Write disabling R1/R2 volumes on page 470.
- **RW Disable R2s** — See Read/write disabling R2 volumes on page 468.
- **Refresh** — See Refreshing R1 or R2 volumes on page 471.
- **Set Mode** — See Setting SRDF mode on page 446.

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

Action	Activate Type	Write Pacing Type	Description
			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
(R)DFA Mode	S = Single-session M = MSC - = N/A

Flag	Status
(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.

Understanding RecoverPoint

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

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 - 5876 on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

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 Enginuity 5875.139.93 or higher on the Symmetrix system.
- This feature is not supported on storage systems running HYPERMAX OS 5977.

To tag/untag volumes at the volume level:

From the Volumes dashboard:

1. Select the storage system.
2. Select **Storage › Volumes** to open the **Volume Dashboard**.
3. In the **All Volumes** panel, expand the type of volume you want to tag/untag.
4. The **Volumes** list view opens.
5. Do one of the following:
   - To tag volumes, select volumes, click [>>], and select **Tag for RecoverPoint**.
   - To untag volumes, select volumes, click [>>], and select **Untag for RecoverPoint.**
6. Click **OK** in the confirmation message.

Untagging from the RecoverPoint tagged volumes view:

1. Select the storage 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 storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. To tag the storage group, select it, click [>>], and select **Tag for RecoverPoint**.
   To untag the storage group, select it, click [>>], and select **Untag for RecoverPoint**.
4. Click **OK** in the confirmation message.

Viewing RecoverPoint sessions

1. Select a storage 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 storage 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 storage system.
2. Select **Data Protection › OpenReplicator › RecoverPoint Sessions** to open the **RecoverPoint Sessions** list view.
3. 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.
   - **Protected Tracks** — Number of protected tracks.

Viewing RecoverPoint tagged volumes

1. Select the storage system.
2. Select **Data Protection › Open Replicator › Recover Point Volumes** to open the **RecoverPoint Volumes** list 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** — Viewing RecoverPoint tagged volume details below
   - **Untag for RecoverPoint** — Tagging and untagging volumes for RecoverPoint on page 518

Viewing RecoverPoint tagged volume details

1. Select the storage system.
2. Select **Data Protection › Open Replicator › Recover Point Volumes** to open the **RecoverPoint Volumes** list 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** — Storage system on which the volume resides.
- **Symmetrix Volume ID** — Storage 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.
- **Compressed Size (GB)** — Size of the compressed volume.
- **Compressed Ratio (%)** — Percentage of volume compressed.
- **Compressed Size Per Pool (GB)** — Size of the compressed pool.
- **Optimized Read Miss** — Cacheless readmiss status. Possible values are:
  - **Off** — Feature is disabled.
  - **System Managed** — Symmetrix system determines the appropriate optimized read miss mode.

The following controls are available:

- **Untag for RecoverPoint** - See Tagging and untagging volumes for RecoverPoint on page 518.

Protecting storage groups using RecoverPoint

This procedure explains how to protect storage groups using RecoverPoint.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 through 5876. RecoverPoint is not supported on storage systems running HYPERMAX OS 5977 or higher.
- The storage group being replicated must be masked to the host.
- The storage group being replicated must not contain any volumes that are already tagged for RecoverPoint.
- Connectivity to the RecoverPoint system/cluster is available.
- RecoverPoint 4.0 is setup and operational. For each cluster in the setup, gatekeepers and repository volumes must be configured in their relevant masking view. Unisphere for VMAX uses a default journal masking view naming convention.
- You must have StorageAdmin permission.
- Depending on the options selected as part of the Protect Storage Group wizard and the existing configuration, some values for some options might populate automatically.

To protect storage groups using RecoverPoint:

1. Select the storage system.
2. Select **Storage › Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.
4. On the **Select Protection Type** page, select **Continuous Protection** , and then **Using RecoverPoint**.
5. Click **Next**.
6. On the **Group Production Settings** page, specify the following information:
   - **RecoverPoint System** — RecoverPoint system.
   - **RecoverPoint Group Name** — Name of the RecoverPoint group.
   - **RecoverPoint Cluster** — RecoverPoint cluster.
   - **Production Name** — Name of the production.
   - **Journal Thin Pool** — Journal thin pool.
   - **Journal Port Group** — Journal port group.
7. Click **Next**.
8. On the **Add Copies** page, add one or more copies. To add a copy, complete the following steps:
   a. Specify the following information:
      – **RecoverPoint Cluster** — RecoverPoint cluster.
      – **Copy Name** — Name of the RecoverPoint copy.
      – **Mode** — Specify whether the mode is Synchronous or Asynchronous.
      – **Symmetrix** — Storage system.
      – **Target Storage Group** — Specify whether the RecoverPoint copy targets a new storage group or an existing group.
   b. Click **Add Copy**.

      The copy is listed in the **Copy Summary** table.
9. Click **Next**.
10. On the **Finish** page, verify your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
11. Do one of the following:
    - 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 851 and Previewing on page 850.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Viewing RecoverPoint volumes

This procedure explains how to view the RecoverPoint volumes for a particular consistency group.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint copies :

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click **View Details** to open its details view.
5. In the **Related Objects** panel, click **Replication Sets** to open the Replication Sets list view.
6. Select the replication set and click **View Details** to open its details view.
7. In the **Related Objects** panel, click **Volumes** to open the **Volumes** list view.

   The following properties display:

   - **Volume Name** — Name of the volume.
   - **Volume Capacity (GB)** — Capacity, in GB, of the volume.
   - **Replication Set** — RecoverPoint replication set.
   - **Copy Name** — RecoverPoint copy.
   - **Storage Type** — Type of storage system.
   - **Array Serial Number** — Serial number.
   - **Vendor** — Vendor of the volume.
   - **Product Name** — Storage product installed.
   - **Product Model** — Model of storage system.

Viewing RecoverPoint clusters

This procedure explains how to view RecoverPoint clusters.

Before you begin:

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

To view RecoverPoint clusters:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Select a RecoverPoint system.
4. Select **View Details**. The **RecoverPoint Clusters** table view displays.

   The following information displays:

   - **RPA Cluster Name** — Name of the cluster.
   - **RPAs** — Number of RecoverPoint appliances.
   - **Management IPv4** — IP address, in IPv4 format. If an IPv6 address is used, this column has the value "N/A".
   - **Mangement IPv6** — IP address, in IPv6 format. If an IPv4 address is used, this column has the value "N/A".
   - **RPA Type** — RecoverPoint appliance type.
   - **Maintenance Mode** — Maintenance mode in use.

Viewing RecoverPoint cluster details

1. Select the storage system.
2. Select **Data Protection › Data Replication and Pools › Recover Point Systems** to open the **RecoverPoint Systems** list view.
3. Select a system and click **View Details** to open the cluster list view.
4. Select a cluster and click **View Details** to open the cluster details view.

   This view allows you to view the cluster details.

   The following properties display:

   - **Cluster Name** — Volume name.
   - **Management IPv4** — IP address, in IPv4 format. If an IPv6 address is used, this column has the value "N/A".
   - **Mangement IPv6** — IP address, in IPv6 format. If an IPv4 address is used, this column has the value "N/A".
   - **Software Serial ID** — Serial ID of the software.

- **RPAs** — Number of RecoverPoint appliances.
- **RPA Type** — RecoverPoint appliance type.
- **Time Zone** — Time zone.
- **Maintenance Mode** — Maintenance mode in use.
- **Internal cluster name** — Internal name of the cluster.

Viewing RecoverPoint splitters

This procedure explains how to view RecoverPoint splitters.

Before you begin:

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

To view RecoverPoint splitters:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Select a RecoverPoint system.
4. Select a system and click **View Details** to open the cluster list view.
5. Select a cluster and click **View Details** to open the cluster details view.
6. In the **Related Objects** pane, click **Splitters** to open the RecoverPoint splitters list view.

   The following information displays:

   - **Splitter Name** — Name of the splitter.
   - **Splitter Array ID** — Array ID of the splitter.
   - **Splitter Array Type** — Array type of the splitter.
   - **Status** — Status of the splitter.
   - **Attached RPA Clusters** — Number of attached clusters.

Viewing RecoverPoint appliances

This procedure explains how to view RecoverPoint appliances.

Before you begin:

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

To view RecoverPoint appliances:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Select a RecoverPoint system.
4. Select a system and click **View Details** to open the cluster list view.
5. Select a cluster and click **View Details** to open the cluster details view.
6. In the **Related Objects** pane, click **RecoverPoint Appliances** to open the RecoverPoint appliances list view.

   The following information displays:

   - **RPA Name** — Name of the RecoverPoint appliance.
   - **Status** — Status of the RecoverPoint appliance.
   - **Private IP** —
   - **Management IPv4** — IP address, in IPv4 format.
   - **Local RPA Fibre Connectivity** —
   - **Remote RPA Connectivity** —

RecoverPoint systems

Discovering RecoverPoint systems

This procedure explains how to discover RecoverPoint systems.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 through 5876. RecoverPoint is not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation you must be a StorageAdmin.

To discover RecoverPoint systems:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Click **Discover RecoverPoint System** to open the **Discover RecoverPoint System** dialog box.
4. In the **Discover RecoverPoint System** dialog box, type the following information:
   - **System Name** — RecoverPoint system name.
   - **System IPv4** — System IP address, in IPv4 format.
   - **System Port** — System port number.
   - **System Username** — System username.
   - **System Password** — System password.
   - **Confirm System Password** — Re-enter system password.
5. Select **OK**.

Deleting RecoverPoint systems

Before you begin:

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

To delete RecoverPoint systems:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Select one or more systems and click **Delete**.
4. Click **OK** in the confirmation message.

Updating RecoverPoint discovery information

This procedure explains how to update the discovery information for a RecoverPoint system.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 through 5876. RecoverPoint is not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation you must be a StorageAdmin.

To update RecoverPoint discovery information:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Select a RecoverPoint system.
4. Click **Update Discover Information** to open the **Update Discover RecoverPoint System** dialog box.
5. In the **Update Discover RecoveryPoint System** dialog box, type the following information:
   - **System Port** — System port number.
   - **System Username** — System username.
   - **System Password** — System password.
   - **Confirm System Password** — Re-enter system password.
6. Click **OK**.

Viewing RecoverPoint systems

This procedure explains how to view previously discovered RecoverPoint systems.

Before you begin:

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

To view RecoverPoint systems:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.

   The RecoverPoint Systems view contains two tabs (**Local Systems** and **All Systems**). The following properties display, depending on the tab:

   Local Systems *tab:*

   - **System Name** — Name of the local system.
   - **IPv4 Address** — IP address of the system.
   - **Clusters** — Number of RPA clusters in the system.
   - **Consistency Groups** — Number of consistency groups associated with the system.
   - **Error Events** — Number of events reported for the system.
   - **Error Alerts** — Number of alerts reported for the system.

   All Systems *tab:*

   - **System Name** — Name of the system.
   - **IP** — IP address of the system
   - **Clusters** — Number of RPA clusters in the system.
   - **Consistency** — Number of consistency groups associated with the system.
   - **Error Events** — Number of events reported for the system.
   - **Error Alerts** — Number of alerts reported for the system.

   The following controls are available on both tabs:

   - **Discover RecoverPoint System** —RecoverPoint systems on page 529
   - **Update Discovery Information** —Updating RecoverPoint discovery information on the previous page
   - **View Details** —Viewing RecoverPoint system details on the next page
   - **Delete** —Deleting RecoverPoint systems on page 529

Viewing RecoverPoint system details

Before you begin:

RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.

To view RecoverPoint system details:

1. Select the storage system.
2. Select **Data Protection › Replication Groups and Pools › RecoverPoint Systems** to open the **RecoverPoint Systems** view.
3. Select the system and click View Details to open its Details view.

   The following properties display:

   - **RPA Cluster Name** — Name of the RPA cluster.
   - **RPAs** — Number of RecoverPoint appliances.
   - **Management IPv4** — IP address, in IPv4 format. If an IPv6 address is used, this column has the value "N/A".
   - **Management IPv6** — IP address, in IPv6 format. If an IPv4 address is used, this column has the value "N/A".
   - **RPA Type** — RecoverPoint appliance type.
   - **Management Mode** — Maintenance mode in use.

RecoverPoint consistency groups

Viewing RecoverPoint consistency groups

This procedure explains how to view the consistency groups used to protect the RecoverPoint volumes.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint consistency groups :

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.

   Use the **Consistency Group** list view to view consistency groups on the selected RecoverPoint system.

   The following properties display:

   - **Consistency Group** — Consistency group name.
   - **Group State** — Consistency group state.
   - **Link States** — Link state.
   - **Source Capacity (GB)** — Source capacity, in GB.
   - **Primary RPA** — Primary RecoverPoint appliance number.
   - **Production Copy** — Production copy name.

   The following controls are available:

   - **View Details** — Viewing RecoverPoint consistency group details
   - **Copies** — RecoverPoint copies
   - **Replication Sets** — RecoverPoint replication sets
   - **Links** — RecoverPoint links

Viewing RecoverPoint consistency group details

This procedure explains how to view the details of a consistency group that is used to protect the RecoverPoint volumes.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint consistency group details:

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a consistency group and click **View Details**to view the properties of that consistency **Consistency Group**.

   Use the **Consistency Group** details view to view consistency group details.

   The following properties display:

   - **Group State** — State of the group.
   - **Group Setting** —
   - **Production Copy** — Name of the production copy.
   - **Copies** — Number of associated copies.
   - **Replication Sets** — Number of associated replication sets.
   - **Active Links** — Number of active links.
   - **Passive Links** — Number of passive links.
   - **Link States** —Lists the states of associated links.
   - **Distributed Group** —
   - **Managed by RecoverPoint** — Indicates if the consistency group is managed by RecoverPoint.
   - **Read Only Replica Volumes** —

RecoverPoint copies

Viewing RecoverPoint copies

This procedure explains how to view the RecoverPoint copies for a particular consistency group.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint copies :

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click **Copies** to open the Copies list view.

   Use the **Copies** list view to view consistency groups on the selected RecoverPoint system.

   The following properties display:

   - **Name** — Name of the copy.
   - **State** — State of the copy. Valid values include **Enabled** and **Suspended**.
   - **Copy Size** — Size of the copy.
   - **Copy Role** — Current role of the copy. Valid values include **Active** and **Replica**.
   - **RTO** —
   - **Journal State** — Indicates the state of the journal. Valid values include **Locked** and **Distributing**.
   - **Journal Size (GB)** — Size of the journal, in GB.

   The following controls are available:

   - **View Details** — Viewing RecoverPoint copy details

Viewing RecoverPoint copy details

This procedure explains how to view the details of a RecoverPoint copy.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint copy details:

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecovertPoint consistency group and click **Copies** to open the **Copies** list view.
5. Select a copy and click **View details** to open its details view. The following properties display for the copy:
   - **Name** — Name of the copy.
   - **State** — State of the copy. Valid values include **Enabled** and **Suspended**
   - **Role** —
   - **Copy Size**
   - **Journal Size**
   - **Journal State**
   - **Journal Volume Name**
   - **Cluster**
   - **RTO (Seconds)**
   - **Journal Size Limit**
   - **AllowDistribOfLargeSnaps**
   - **AllowSymmWithOneRPA**
   - **ActivePrimaryRPA**
   - **FastForwardBound**
   - **NumCopySplitters**
   - **NumCopyVolumes**
   - **NumJournalVolumes**
   - **PhoenixDevices**
   - **TspWritesCleared**

- **UserSnapshot**
- **Production Copy** — Production copy.
- **Volumes** — Number of associated volumes.
- **Copy Capacity (GB)** — Capacity of the copy, in GB.
- **Journal Size (GB)** — Size of the journal, in GB.

RecoverPoint replication sets

Viewing RecoverPoint replication sets

This procedure explains how to view the RecoverPoint replication sets for a particular consistency group.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint replication sets:

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click **Replication Sets** to open the **Replication Sets** list view.

   Use the **Replication Sets** list view to view replication sets associated with the selected consistency group.

   The following properties display:

   - **Name** — Name of the replication set.
   - **Capacity (GB)** — Source capacity, in GB.
   - **Production Volume Capacity (GB)** — Production volume capacity, in GB.
   - **Volumes** — Number of associated volumes.

   The following controls are available:

   - **View Details** — Viewing RecoverPoint replication set details

Viewing RecoverPoint replication set details

This procedure explains how to view the details of a RecoverPoint replication set.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint set details:

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click **Replication Sets** to open the **Replication Sets** list view.
5. Select a replication set.
6. Click **View details**. The following properties display for the replication set:
   - **Name** — Name of the replication set.
   - **Volumes** — Number of associated volumes.
   - **Volume Name** — Number of associated volumes.
   - **Production Volume Capacity (GB)** — Production volume capacity, in GB.
   - **Capacity (GB)** — Source capacity, in GB.

RecoverPoint links

Viewing RecoverPoint links

This procedure explains how to view the RecoverPoint links for a particular consistency group.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint links :

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click **Links** to open the **Links** list view.

   Use the **Links** list view to view the links associated with the selected consistency group.

   The following properties display:

   - **Name** — Name of the RecoverPoint link.
   - **Transfer Enabled** — Indicates whether transfer is enabled for this RecoverPoint link.
   - **Link State** — Indicates whether the link state is active or paused.
   - **Local** — Indicates whether or not the link is local.
   - **Protection Mode** —

   The following controls are available:

   - **View Details** — Viewing RecoverPoint link details

Viewing RecoverPoint link details

This procedure explains how to view the details of a RecoverPoint link.

Before you begin:

- RecoverPoint operations on Unisphere for VMAX require Enginuity 5875.139.93 or higher on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- To perform this operation, you must be a monitor or higher.

To view RecoverPoint link details:

1. Select the storage system.
2. Select **Data Protection › RecoverPoint** to open the **RecoverPoint** view.
3. Select a RecoverPoint system and click **View Details**to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click **Links** to open the **Links** list view.
5. Select a link.
6. Click **View details**. The following properties display for the link:
   - **Name** — Name of the RecoverPoint link.
   - **Transfer State** — Indicates if the transfer state is enabled.
   - **Link State** — Indicates if the link state is active or paused.
   - **Local** — Indicates if the link is local.
   - **RPO (seconds)** — Indicates the RPO, in seconds.
   - **First Copy** —
   - **Second Copy** —
   - **Protection Mode** —
   - **Replication Over WAN** — Indicates if replication over WAN is supported.
   - **WAN Compression** — Specifies what WAN compression, if any, is being used.
   - **Bandwidth Limit** —
   - **Deduplication** — Specifies if deduplication is enabled.
   - **Perform Long Initialization** —
   - **Snapshot Granularity**—

Open Replicator

Creating Open Replicator copy sessions

Before you begin:

When the ORS control volumes in on a storage system running HYPERMAX OS 5977 or higher, the following session options cannot be used.

- Push
- Differential
- Precopy

There are many rules and limitations for running Open Replicator sessions. Refer to the *Solutions Enabler Migration CLI Product Guide* before creating a session. For quick reference refer to Open Replicator session options on page 548.

To create a copy session:

1. Select the storage 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.

    This feature is not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.

11. Select the Open Replicator session options (page 548) 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 storage 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 (director 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 page 541.
- **Create FLM Session** — See Creating a FLM session on page 555. (Not supported for volumes on storage systems running HYPERMAX OS 5977)
- **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 storage 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 548 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.
- Restore operations are not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977

To restore a session:

1. Select the storage 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 548 for session control options.
5. Click **OK**.

Renaming Open Replicator sessions

Before you begin:

Renaming operations are not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.

To rename Open Replicator sessions:

1. Select the storage 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

Before you begin:

Removing Open Replicator sessions is not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.

To remove sessions:

1. Select the storage 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.

Before you begin:

Setting background copy mode to precopy is not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.

To set background copy mode:

1. Select the storage 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 548 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 storage 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 548.
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 storage 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 548 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.

Before you begin:

This feature is not supported on storage systems running HYPERMAX OS 5977 or higher.

Setting Open Replicator session pace:

1. Select the storage 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 storage 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 storage 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 548 for session control options.
5. Click **OK**.

Viewing Open Replicator sessions

1. Select the storage 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 next 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 543.
- **Terminate** — See Terminating Open Replicator sessions on the previous page.
- **Remove** — See Removing Open Replicator sessions on page 544.
- **Front End Zero off** — See Setting Open Replicator session donor update off on page 545.
- **Donor Update Off** — See Setting Open Replicator session donor update off on page 545.
- **Rename** — See Renaming Open Replicator sessions on page 544.
- **Remove** — See Removing Open Replicator sessions on page 544.
- **Restore** — See Restoring Open Replicator sessions on page 543.
- **Set Pace** — See Setting Open Replicator session pace on page 545.
- **Set Mode** — See Setting Open Replicator session background copy mode on page 544.

Viewing Open Replicator session details

1. Select the storage 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

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

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.

Session Option	Used with Command	Description
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). This option is not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.
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.
Front-End Zero Detection	Create	Enables front end zero detection for thin control volumes in the session. Front end zero detection looks for incoming zero patterns from the remote volume, and instead of writing the incoming data of all zeros to the thin control volume, the group on the thin volume is de-allocated.
Hot	Create	Hot copying allows the control device to be read/write online to the host while the copy operation is in progress. All directors that have the local devices mapped are required to participate in the session. A hot copy session cannot be created unless all directors can discover the

Session Option	Used with Command	Description
		remote device.
Nocopy	Activate	Temporarily stops the background copying for a session by changing the state to CopyOnAccess or CopyOnWrite from CopyInProg.
Pull	Create	A pull operation copies data to the control device from the remote device.
Push	Create	A push operation copies data from the control volume to the remote volume. This option is not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.
Precopy	Create Recreate	For hot push sessions only, begins immediately copying data in the background before the session is activated. This option is not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977
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.
Z	X =Enabled

Flag	Status
Front-end zero detection	. = Disabled
*	Failed session can be reactivated.

Understanding Federated Live Migration

Federated Live Migration (FLM) allows you to move data from older storage into a VMAX Series system nondisruptively. The host application cutover to use the 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 Federated Live Migration is not supported on storage systems running HYPERMAX OS 5977. For information on supported operating systems, file systems, and logical volume managers, refer to the *Federated Live Migration Simple Support Matrix*.

An example FLM configuration including the network, storage systems, application hosts, and Unisphere for VMAX host is shown below:



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 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 on the facing page

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 volumes to maintain data availability. For details on properly configuring your host environment, refer to *Federated Live Migration Technical Overview Technical Notes*.

Procedure overview:

There are three primary tasks involved with setting up an running a Federate Live Migration:

- Readying the environment
- Creating the migration session and masking view
- Migrating the data

Ready the environment for a migration:

1. Verify that the application is running on the source storage system and that the storage system is zoned to the application and SYMAPI servers, as shown next:



2. Create/modify zoning to support connectivity of the target storage system to the:
   - Application Server (application access to the target storage system)
   - SYMAPI server (control of the target storage system)
   - Source storage system (application data transfer)

   As shown next:



Create 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 the facing 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 546.
2. Continue with the remaining steps in this procedure.

Migrate the data

To migrate the data:

1. Activate the Federated Live Migration session.

2. Once the migration completes (i.e., all the data has been transferred to the target storage system), terminate the Open Replicator session, as described in Terminating Open Replicator sessions on page 546, and remove the zoning and masking of the source storage system 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 *Solutions Enabler Migration CLI Product Guide*.

Note The host may run indefinitely with federated identity on the new storage system volumes; however, it is recommended 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; storage system volumes may remain federated indefinitely and there is no requirement to unspoof at any time.

The end result is a source storage 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 Data Migration CLI Product Guide* before creating a session.

Before you begin:

This feature is not supported on storage systems running HYPERMAX OS 5977.

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 542.
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**: Type 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**: Type 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 management

This chapter contains the following:

Common Tasks 560
Monitor view 566
Analyze view 578
Charts view 613
Plan view 620
Reports view 623
Settings 629
Metrics 642

Common Tasks

This chapter contains the following:

Exporting Performance settings

You can import the following items:

- Performance metrics (KPIs) definitions with their default (not instance-specific) settings.
- Global alert definitions.
- User template dashboards.

For more information about exporting and importing Performance settings, see About exporting and importing Performance settings on the facing page.

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, ensure that you also provide the password.

To export Performance settings:

1. Click **Performance**.
2. Under **Common Tasks**, click **Export Settings** to open the **Export Settings** dialog box.
3. Select the settings to export: **Metric and Alert Definitions**, **User Templates**, or both.
4. Enter and confirm a file password.

   The settings file cannot be imported to another Unisphere environment without supplying the password set in this step.
5. Click **OK**. The **Save Settings** confirmation appears, click **OK**.
6. In the **Save File** dialog box select the location to which you want to save the settings file

Importing Performance settings

You can import the following items:

- Performance metrics (KPIs) definitions with their default (not instance-specific) settings.
- Global alert definitions.
- User template dashboards.

For more information about exporting and importing Performance settings, see About exporting and importing Performance settings on the facing page.

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, ensure that you also provide the password.

To import Performance settings:

1. Click **Performance › Import Settings** to open the **Import Settings** dialog box.
2. Select the settings to import: **Metrics and Alert Definitions**, **User Templates**, or both.
3. Click **Browse** to select the settings file to import.
4. Enter the file password (which was set during the export process). Click **OK**.

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. Click **Performance**.
2. Under **Common Tasks**, click **Export PV Settings** to open the **Export Performance Viewer Settings** dialog box.
3. Select the settings to export: **SYMAPI Database**, **Settings**, or both.

   If you select **SYMAPI Database**, specify a file name for the BIN file.

   If you select Settings, specify a file name for the ZIP file. In addition, select the types of settings you want to export. The following settings are available:

   - Thresholds (default and custom)
   - Reports
   - User Defined Dashboards
   - Templates
4. 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: `install_dir\SMAS\backup\SPA\SE`

   Settings: `install_dir\SMAS\backup\SPA\settings`

About exporting and importing Performance settings

In large environments, it can be very time consuming to setup and configure multiple instances. storage systems may have very similar configuration and implementations across the environment. To assist with the implementation process, the following settings can be exported and imported:

- Performance metrics (KPIs) definitions with their default (not instance-specific) settings.

- Global alert definitions.
- User template dashboards.

The settings are saved to an XML file. You can select the client-side directory in which the file is saved.

When importing a metrics definition file, it replaces any existing metrics definition file.

When importing a user template dashboard, any existing templates remain intact; Only non-conflicting templates are imported.

Saving dashboards and charts

When viewing a dashboard or chart, you can export it to a PDF file or save it as an image file (JPG or PNG file format) . When exporting to PDF, you can export the chart/dashboard only, or the dashboard/chart with the data also included in a table.

If multiple dashboards are displayed on different tabs, exporting to a PDF saves each chart or dashboard. Saving a dashboard or chart as an image file saves only the currently displayed dashboard/chart.

You can use the **Save As › To Clipboard** functionality open a JPG image of the current dashboard in a new browser tab. From there, you can copy it to the clipboard and paste it to another application, for example, into a presentation.

Note This functionality is available only when using Firefox or Chrome web browsers.

To save a dashboard to an image file:

1. Open the dashboard or chart that you want to save.
2. Click **Save As**.
3. Select the image format you want to use. The following choices are available:
   - JPG
   - PNG
4. In the **Save As** window, navigate to the folder to which you want to save the image, and change the file name, if required.
5. Click **Save**.

To export a dashboard to PDF:

1. From the **Monitor** view or the **Plan** view, open the dashboard you want to save.
2. Click **Save As**.
3. Select the information you want to export to PDF. The following choices are available:
   - PDF - Chart
   - PDF - Charts & Data
4. In the **Save Reports Results** window, click **OK**.
5. In the **Save As** window, navigate to the folder to which you want to save the image, and change the file name, if required.
6. Click **Save**.

To open a dashboard image in a new browser tab:

1. Using Firefox or Chrome, open the dashboard or chart that you want to save.
2. Click **Save As › To Clipboard**.
3. A JPG image of the current dashboard opens in a new browser tab.

Filtering performance data

Unisphere for VMAX includes different ways to filter performance data depending on the situation:

- Filter performance charts and heatmaps by name. This allows you filter for specific objects by name. For example, the name of a storage group.
- Filter performance charts and heatmaps by ID. This allows you to filter lists according to an ID of a particular object. For example, the ID of a volume.
- Filter performance charts and heatmaps by object. This allows you to filter lists according to an instance of a particular object. For example, a particular storage group or database.
- Filter analyze views by name.

To filter performance charts and heatmaps:

1. Click the filter icon to open the filter dialog box.
2. Depending on the type of filter, some of the following options are available:
   - Filter by name. To do this, in the **Name** field, type all or part of the object's name.
   - Filter by instance. To do this, complete the following steps:
     a. From the **Context** menu, select an object type. The **Instances** list updates to show the objects of that type.
     b. Select the object from the **Instances** list. You can filter the **Instances** list by typing all or part of the object's name in the **Instances** field.
   - Filter by ID. To do this, complete the following steps:
     a. From the **Context** menu, select **By ID**.
     b. In the **ID** field, type all or part of the ID.
3. Click **OK**.

To filter analyze views:

1. Click the filter icon to expand the **Filtered by** field.
2. In the **Filtered by** field, type all or part of the name.

   To clear the filter, click **Reset**.

Changing the time range

Use the options in the Custom Time Selection dialog to specify a custom time range for performance reports or to specify a data collection time range for an analyze view.

In the **Analyze** tab, there are the following default data collection time ranges:

- Real Time Analysis — The previous 1 hour.
- Root Cause Analysis — The previous 4 hours.
- Trending & Planning — The previous 7 days.

For the **Real Time Analysis** view, you can load an existing real-time trace file. After a trace file is specified, the time range is displayed to the left of the time selection menu.

For the **Root Cause Analysis** and **Trending & Planning** views, you can customize the time range by specifying a custom time range in the **Custom Time Selection** window.

When specifying a custom Root Cause Analysis time range, you can define a maximum range of 24 hours, specified at five-minute intervals.

When specifying a custom Trending & Planning time range, you can define a minimum range of 24 hours, specified at one-hour intervals.

After a custom time range is specified, the time range is displayed to the left of the time selection menu.

To change the Real Time Analysis time range:

1. From the time selection menu, which, by default displays 'Last Hour', select one of the following options:
   - Select **Last Hour** to view the last hour of data collected.
   - Select **Trace** and choose a saved trace file from the list.
     This option is only available if a trace file is available for the selected storage system.

To change the Root Cause Analysis time range:

1. From the time selection menu, which by default displays 'Last 4 Hours', select one of the following options:
   - Select **Last Hour** to view the last hour of data collected.
   - Select **Last 2 Hours** to view the last two hours of data collected.
   - Select **Last 4 Hours** to view the last four hours of data collected.
   - Select **Last 8 Hours** to view the last eight hours of data collected.
   - Select **Last 24 Hours** to view the last 24 hours of data collected.
   - Select **Custom** to make a custom time selection. To do this, complete the following steps:
     1. In the **Custom Time Selection** window, specify a start time and end time.
     2. Click **OK** .

To change the Trending & Planning time range:

1. From the time selection menu, which by default displays 'Last 24 Hours', select one of the following options:
   - Select **Last Hour** to view the last hour of data collected.
   - Select **Last 2 Hours** to view the last two hours of data collected.
   - Select **Last 4 Hours** to view the last four hours of data collected.
   - Select **Last 8 Hours** to view the last eight hours of data collected.
   - Select **Last 24 Hours** to view the last 24 hours of data collected.

- Select **Custom** to make a custom time selection. To do this, complete the following steps:
    1. In the **Custom Time Selection** window, specify a start time and end time.
    2. Click **OK** .

To set a custom report time range:

1. In the **Custom Time Selection** window, specify a start time and end time.
2. Click **OK** .

Monitor view

This chapter contains the following:

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 User Dashboards folder. You can use folders 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 Views**.

Charts and heatmaps that you create can be saved, edited, and deleted.

The dashboards under **EMC Views** are created dynamically and provide the following benefits:

- Easy visualization.
- Event correlation between application response time and hardware component utilization taken at a particular point in time.
- Fast problem recognition.

You can save an EMC dashboard as a template.

The heatmap under **EMC Views** displays hardware instances represented as colored squares, with the color indicating utilization levels. Clicking on any square displays Utilization and Queue Depth information for that hardware instance.

The specific information displayed in a heatmap depends on the model of the selected storage system and whether any filters have been applied.

The specific information displayed in a Utilization summary dashboard view depends on the selected Symmetrix ID and Category. Unless the selected Category is Array, the information is filtered for the selected category (for example, Database or Storage Group).

To manage dashboards:

Creating a dashboard or template folder on page 569

Creating a dashboard with charts on page 569

Copying a dashboard on page 569

Saving a dashboard as a template on page 571

Viewing dashboards on page 568

Deleting a dashboard on page 570

Saving dashboards and charts

To manage EMC dashboards:

Managing EMC Views on page 572

To manage EMC summary dashboards:

Using EMC summary dashboards on page 573
Scheduling a report from the dashboard on page 572

To manage projection dashboards:

Array projection dashboard
Thin pool projection dashboard

Viewing dashboards

Before you begin:

Read Managing dashboards on page 566.

To view dashboards:

1. Select **Performance › Monitor**. By default, the **Dashboards** view is displayed.

   The **Dashboards** 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 Views** — This folder contains a predefined dashboard, summary and heatmap. The dashboards can be modified dynamically, but cannot be deleted. They can be modified and saved as a template.

   When viewing a dashboard, the right section of the view displays the selected dashboard. The bottom of the view block provides dynamic configuration controls, when applicable.

   Depending on the selected dashboard, some of the following controls are available:

   - Save as — A drop-down menu containing some or all of the following options:
     - Dashboard — Creating a dashboard with charts on the facing page
     - Template — Saving a dashboard as a template on page 571
     - JPG — Saving dashboards and charts on page 562
     - PNG — Saving dashboards and charts on page 562
     - PDF Charts — Saving dashboards and charts on page 562
   - **Run Report Now** — Running a report from the dashboard on page 571
   - **Schedule Report** — Scheduling a report from the dashboard on page 572
   - **Edit** — Editing a dashboard on page 570
   - **Save** — Saving dashboard changes on page 571
   - **Copy** — Copying a dashboard on the facing page
   - **Delete** — Deleting a dashboard on page 570
   - **Navigate to Details View** — Navigating to the Details view on page 1
   - **Navigate to Analyze** — Navigating from a heatmap to an analyze view on page 572

Creating a dashboard or template folder

You can create a new folder for dashboards or templates manually. Alternatively, you can create a new folder as part of the process when creating a new dashboard or saving a new template.

To create a dashboard or template folder:

1. Select **Performance** › **Monitor** to open the **Monitor** view.
2. Right-click on the **User Dashboards** or **User Templates** folder in the dashboards tree. Select **New Folder**.
3. Enter a folder name and click **OK**.

   The folder displays in the **User Dashboards** or **User Templates** section of the dashboards tree.

Creating a dashboard with charts

To create a dashboard with charts:

1. Select component(s) from the table in Analyze view, and create one or more charts.
2. Click **Save As Dashboard**. The **Create Dashboard** dialog box displays.
3. Specify a dashboard name and folder.
4. *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.
5. Click **OK** to save the dashboard.

Copying a dashboard

You can copy an existing dashboard and change the name and a few characteristics to make a new dashboard.

To copy a dashboard:

1. Select the dashboard to copy and click **Copy** to open the **Copy Dashboard** or **Copy Template** dialog box, depending on the dashboard type.
2. Type a **Name** for the copied dashboard.
3. Select a **Folder** for the dashboard location, or select **Create new folder**.
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.
5. *Template Dashboards only*: Select the **Instances** for the new dashboard.
6. *Optional*: Select a chart from the list of charts and use **Delete** or **Edit** to customize the copied dashboard.
7. Click **OK**.

Editing a template dashboard

To edit a template dashboard:

1. Select **Performance › Monitor** to view the list of user dashboards and user template dashboards.
2. In the dashboards tree, select the user template dashboard you want to edit and click **Edit**. The **Edit Template** dialog displays.
3. *Optional*: Change the dashboard **Name**.
4. *Optional*: Change the dashboard **Folder**.
5. *Optional*: Change the value in **Instances**.
6. *Optional*: Select an item from the list of charts. Perform one or more of the following actions:
   - Click **Edit** to open the **Edit Snapshot Chart** dialog box.
     Change the chart **Title**, **Symmetrix ID**, **Chart Style**, **Data Format**, or **Time Range** and click **OK**.
     Note that depending on the dashboard type, not all options are available.
   - Click **Delete** to delete the chart from the list. Deleting a chart requires a confirmation **OK**.
   - Click **Add** to add a chart to the list.
7. Click **OK**.

Editing a 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 list of charts and click **Edit** to open the **Edit Snapshot Chart** dialog box.

   Change the chart **Title**, **Symmetrix ID**, **Chart Style**, **Type**, **Data Format**, or **Time Range** and click **OK**.
5. Click **OK**.

Deleting a dashboard

Before you begin:

You can delete only user-defined dashboards. Predefined dashboards cannot be removed.

To delete a dashboard:

1. Select **Performance › Monitor** to open Monitor view.
2. Select the dashboard you want to delete from the dashboard navigation tree and click **Delete**.
3. Click **OK** on the delete confirmation message.

Running a report from the dashboard

The performance dashboard reports are saved in PDF format. Dashboard reports show the charts and summary reports as they appear in the dashboard view. To create reports for other data formats (XML, CSV, and BTP), and to manage any scheduled reports, click **Performance › Reports**.

To run a report from the dashboard:

1. Click **Performance › Monitor** to open the **Monitor** view.
2. Select or create a system dashboard.
3. Do one of the following:
   - Click **Save As › PDF - Charts**.
   - Click **Save As › PDF - Charts & Data**.
4. Click **OK** in the **Save Report Results** dialog to save the query results.
5. Type a file name in the **Save File** dialog box.

   The report (in PDF format) is saved to the specified location.

   Reports that are run by clicking **Run Report Now** do not appear in the **Reports** list.

Saving a dashboard as a template

Selected EMC-defined dashboards can be adjusted and saved as a template for your own dashboard. Depending

To save a dashboard as a template:

1. Select **Performance**, and then **Monitor** to open the **Monitor** view.
2. Select one of the **EMC Dashboards**.
3. Set the dashboard characteristics (Type, Time Range) at the top of the dashboard.
4. Click **Save as › Template** to open the **Save Template** dialog box.
5. Enter a dashboard **Name**, select a **Folder**, and select the **Instances** to monitor.
6. *Optional*: If you have charts saved from Analyze view, click **Add** to add them to the dashboard.
7. 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. Click **Performance › Monitor** to view the list of dashboards.
2. In the dashboards tree, select the dashboard you want to change.
3. Update the chart properties, depending on the options available.

For more information, see Customizing the Charts view on page 618 and Customizing charts on page 616

4. When finished, click **Save**.

Managing EMC Views

To view and manage EMC Views:

1. Select **Performance** › **Monitor** to open the Monitor view. The **Dashboards** view is displayed by default.
2. Select **Dashboards**, **Summary**, or **Heatmap**from the **EMC Views** folder.
3. Make your selections from the fields across the top of the dashboard.

   The charts update dynamically based on your selections.

See Using EMC summary dashboards on the facing page. provides performance information for the specified storage system component.

Navigating from a heatmap to an analyze view

You can navigate from an EMC heatmap object to an analyze view.

To navigate to an analyze view from a heatmap:

1. Select **Performance** › **Monitor** to open the **Monitor** view.
2. Select **EMC Views** › **Heatmap**.
3. *Optional*: Select the **Type** (**Real Time**, **Root Cause Analysis**, or **Trending & Planning**) and the **Time Range** from the drop-down options above the heatmap.
4. Double-click on any object in the heatmap to open its corresponding analyze view.
5. Click **Navigate to Analyze**.

   The **Root Cause Analysis**view opens to the selected object's location.

Scheduling a report from the dashboard

The performance dashboard reports are saved in .pdf format. When you schedule a report from the Monitor view, it will run continuously on the time/days selected.

To schedule a report from the dashboard:

1. Select **Performance** › **Monitor** to open the **Monitor** view.
2. Select or create a system dashboard.
3. Click **Schedule Report** to open the **Create Report** dialog box.
4. In the **General** tab, enter a report **Name**.
5. *Optional*: Enter the report description.
6. In the **Schedule** tab, select **Schedule**.
7. Set **First Run Time**, the time at which the report will run.
8. Select the days on which the report will run.

9. From the **Retention (Days)** menu, select the number of days for which you want the schedule to apply.
10. *Optional*: In the **Email** tab, configure email notifications.
11. Click **OK**.

A confirmation message displays. Any changes to this scheduled report are made in **Performance › Reports**.

Scheduled reports save to this default location: `install_dir\SMAS\report\SPA\querydata\`

Using EMC summary dashboards

Before you begin:

The EMC summary dashboards in **Performance › Monitor › EMC Views › Summary** require **Administrator**, **Monitor**, or **Performance Monitor** permission.

Summary dashboards available:

The following summary dashboards are available in the **Performance › Monitor › EMC Views › Summary** view:

- **BE Director (DA)** — BE Directors (DA) Summary
- **Disk Group** — Disks Summary
- **Disk Technology** — Disks by Technology Summary
- **External Director** — BE Directors (DX) summary
- **External Disk Group** — External Disks Summary
- **FE Director** — FE Directors Summary
- **Port Group** — Port Groups Summary
- **RDF Director** — RDF Directors Summary
- **Storage Group** — Storage Groups Summary

You can save summary dashboards in a PDF file by selecting **Save As › PDF - Charts**, and specifying a file name and directory location.

Summaries can be scheduled by selecting **Schedule Report**. Scheduled reports are listed in the **Performance › Reports** page. When a scheduled report completes, it saves to the default location: `install_dir\SMAS\report\SPA\querydata`.

Arrays summary

The array summary can be scheduled.

The following metrics are available for all storage systems:

- **Host IOs/sec** — The number of host IO operations performed each second by all volumes, including writes and random and sequential reads.
- **Host MBs/sec** — The number of host MBs read by all of the volumes each second.
- **BE Reqs/sec** — The number of read or write requests between the cache and the director.
- **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 is reuse by another address.
- **WP Count** — The number of system cache slots that are write pending.
- **% Cache WP** — The percent of cache that is busy with write pendings.

See Using EMC summary dashboards on the previous page.

See Running a report from the dashboard on page 571.

See Scheduling a report from the dashboard on page 572.

BE Directors (DA) Summary

The following metrics are available for all back-end directors:

- **ID** — The ID of the director.
- **% Busy** — The percent of time the director is busy.
- **% Non-IO Busy** — The percent of time the director is busy with non-IOs.
- **IOs/sec** — The number of IO 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 throughput of host MBs read (per second) by the director.
- **MBs Written/sec** — The average throughput of host MBs written (per second) by the director.

See Using EMC summary dashboards on the previous page.

See Running a report from the dashboard on page 571.

See Scheduling a report from the dashboard on page 572.

Databases Summary

The following metrics are available for all databases:

- **ID** — The ID of the database.
- **Host IOs/sec** — The number of host IO operations performed each second by the database.
- **Response Time (ms)** — The average time that it took the storage system to respond to IO for this database.

- **Host Reads/sec** — The number of host read operations performed each second by the database.
- **Read Response Time (ms)** — The average time that it took the storage system to serve one read IO for database.
- **Host Writes/sec** — The number of host write operations performed each second by the database.
- **Write Response Time (ms)** — The average time that it took the storage group to serve one write IO for this group.
- **% Hit** — The percent of IO operations performed by the database that were immediately satisfied by cache.
- **Allocated Capacity (GB)** — The allocated capacity of the database.
- **Total Capacity (GB)** — The total capacity of the database.

See Using EMC summary dashboards on page 573.
See Running a report from the dashboard on page 571.
See Scheduling a report from the dashboard on page 572.

Disks Summary

The following metrics are available for each disk group:

- **ID** — The ID of the disk.
- **% Busy** — The percent of time that the disk is busy serving IOs.
- **Avg Queue Depth** — Calculated value: Accumulated queue depth / Total SCSI command per sec.
- **Reads/sec** — The number of host reads per second for the disk.
- **Read RT (ms)** — The average time it took the disk to serve one read command.
- **Writes/sec** — The number of host writes per second for the disk.
- **Write RT (ms)** — The average time it took the disk to serve one write command.
- **Total Capacity (GB)** — The total capacity of the disk (GBs).
- **% Used Capacity** — The percent of the total capacity that is used (MBs).

See Using EMC summary dashboards on page 573.
See Running a report from the dashboard on page 571.
See Scheduling a report from the dashboard on page 572.

Disks by Technology Summary

The following metrics are available for each disk technology:

- **ID** — The ID of the disk.

- **% Busy** — The percent of time that the disk is busy serving IOs.
- **Avg Queue Depth** — Calculated value: Accumulated queue depth / Total SCSI command per sec.
- **Reads/sec** — The number of host reads per second for the disk.
- **Read RT (ms)** — The average time it took the disk to serve one read command.
- **Writes/sec** — The number of host writes per second for the disk.
- **Write RT (ms)** — The average time it took the disk to serve one write command.
- **Total Capacity (GB)** — The total capacity of the disk (GBs).
- **% Used Capacity** — The percent of the total capacity that is used (MBs).

See Using EMC summary dashboards on page 573.

See Running a report from the dashboard on page 571.

See Scheduling a report from the dashboard on page 572.

FE Directors Summary

The following metrics are available for all front-end directors:

- **ID** — The ID of the director.
- **% Busy** — The percent of time the directory is busy.
- **Host IOs/sec** — The number of host IOs transferred through the director each second.
- **Host MBs/sec** — The throughput of the director each 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 RT (ms)** — The calculated average response time for reads.
- **Write RT (ms)** — The calculated average response time for writes.

See Using EMC summary dashboards on page 573.

See Running a report from the dashboard on page 571.

See Scheduling a report from the dashboard on page 572.

Storage Groups Summary

The following metrics are available for all storage groups:

- **ID** — The ID of the storage group.
- **Host IOs/sec** — The number of host IO operations performed each second by the storage group.
- **Response Time (ms)** — The average time that it took the storage system to respond to IO for this group.

- **Host Reads/sec** — The number of host read operations performed each second by the group.
- **Read RT (ms)** — The average time that it took the storage group to serve one read IO for this group.
- **Host Writes/sec** — The number of host write operations performed each second by the group.
- **Write RT (ms)** — The average time that it took the storage group to serve one write IO for this group.
- **% Hit** — The percent of IO operations performed by the group that were immediately satisfied by cache.
- **Allocated Capacity (GB)** — The allocated capacity of the storage group.
- **Total Capacity (GB)** — The total capacity of the storage group.

See Using EMC summary dashboards on page 573.

See Running a report from the dashboard on page 571.

Scheduling a report from the dashboard

Analyze view

This chapter contains the following:

Navigating in the Analyze tab

In the **Analyze** tab, you can view storage system data for various collection ranges.

You can use the **Real Time Analysis** and **Root Cause Analysis** views to debug and find causes for performance problems. You can use the **Trending & Planning** view to view historical data, showing long-term trends, to plan for future needs.

- **Real Time Analysis** —The default range is the previous 1 hour. **Real Time Analysis** view collects data between 2 and 5 seconds for a limited group of metrics. The data is available for the previous hour.
- **Root Cause Analysis** —The default range is the previous 4 hours. **Root Cause Analysis** view collects data every 5 minutes for root cause analysis. The data is available for the previous 7 days.
- **Trending & Planning** —The default range is the previous 7 days. **Trending & Planning** view collects data in 15 minute intervals for trending and planning. The data is available for the previous year.

You can use Real Time Analysis and Root Cause Analysis information to debug and find the cause of performance problems. Trending & Planning information shows long-term trends you can use to plan for future requirements.

Each collection range has a default period of time for which data is collected but other ranges are available according to the collection range selected.

Real Time Analysis time ranges:

- Last Hour (default)
- Trace (if available for the selected storage system)

Root Cause Analysis time ranges:

- Last Hour
- Last 2 Hours
- Last 4 Hours (default)
- Last 8 Hours
- Last 24 Hours
- Custom
  This is a user-specified custom time range.

Trending & Planning time ranges:

- Last 24 Hours
- Last Week (default)
- Last 6 Months
- Last 12 Months
- Custom
  This is a user-specified custom time range.

In **Analyze** view you can do the following:

Viewing information in the Analyze view (page 579).

Changing the time range (page 563).

Creating a dashboard from an Analyze view (page 579).

Viewing information in the Analyze view

The **Analyze** view provides three views into your storage data:

- **Real Time Analysis** view
- **Root Cause Analysis** view
- **Trending & Planning** view

Licensed and registered storage systems are displayed in a table format. To see more information about a storage system, click the ID of the storage system. More detailed data, displayed according to category, is presented. You can navigate between the different categories of information for that storage system by selecting a category from the drop-down menu in the breadcrumb trail. Categories for which there is no information available for that storage system are not displayed.

The detailed information displayed for each storage system depends on the configuration and version of that storage system.

To return to the top-level, table view of storage systems, click the storage system ID in the breadcrumb trail.

Creating a dashboard from an Analyze view

You can create a dashboard from the information currently displayed in an **Analyze** view.

Before you begin:

You require access to a visible, registered storage system.

To create a dashboard from an Analyze view:

1. Select component(s) from the table in **Analyze** view, and create one or more charts.
2. Click **Save As › Dashboard**. The **Create Dashboard** dialog box displays.
3. Specify a dashboard name and folder.
4. *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 storage system heat map is the default until you specify a default dashboard.
5. Click **OK** to save the dashboard.

Creating a template dashboard from an Analyze view

To create a dashboard template from an Analyze view:

1. Select component(s) from the table in **Analyze** view, and create one or more charts.
2. Click **Save as Template**. 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.

Viewing a Real Time trace

You can view saved traces of Real Time data to troubleshoot or diagnose performance issues.

Before you begin:

You require one or more saved data traces.

To view a Real Time trace:

1. Select **Performance › Monitor** Click **Analyze** to open the **Analyze** view.
2. Select **Real Time Analysis**.
3. Select the storage system.
4. From the time selection drop-down menu, select **Trace**. The **Trace Selection** dialog box displays.
5. Select a real time trace file.
6. Click **OK**.

When the trace displays in the table, you can create charts for monitoring that time range.

Real Time Analysis views

You can use the Real Time Analysis views to debug and find causes for performance problems.

To open a Real Time Analysis view:

1. Click **Performance › Analyze**.
2. Click **Real Time Analysis**.
3. Click the **Symmetrix ID** of the storage system for which you want to view an analyze view.
4. Select a category from the drop-down menu in the breadcrumb trail to see a view with more information about that category.

This chapter contains information about the following views:

Symmetrix systems view (Real Time Analysis)

The **Real Time Analysis** view provides high-level data about all storage systems. The specific information displayed depends on the models of the selected storage systems. Some or all of the following information is displayed:

- **Host IOs/sec** — The IOs per second for each front end director.
- **Host MBs/sec** — The MBs per second for each front end director.
- **FE 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.

FE Director view (Real Time Analysis)

The **Real Time Analysis** view provides detailed data about front-end directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the director.
- **% Busy** — The percent of time the director 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.

BE Director (DA) view (Real Time Analysis)

The **Real Time Analysis** view provides detailed data about back-end directors (DA). The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 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.

External Director view (Real Time Analysis)

The **Real Time Analysis** view provides detailed data about external directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the external director.
- **% Busy** — The percent of time that the director is busy.
- **IOs/sec** — The number of IOs (per second) handled by the director.
- **Reqs/sec** — The average number of requests performed by the director per second.
- **Read Reqs/sec** — The average number of read requests performed by the director per second.
- **Write Reqs/sec** — The average number of write requests performed by the director per second.

RDF Director view (Real Time Analysis)

The **Real Time Analysis** view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the director.
- **% Busy** — The percent of time the director 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.

Root Cause Analysis views

You can use the Root Cause Analysis views to debug and find causes for performance problems.

To open a Root Cause Analysis view:

1. Click **Performance › Analyze**.
2. Click **Root Cause Analysis**.
3. Click the **Symmetrix ID** of the storage system for which you want to view an analyze view.
4. Select a category from the drop-down menu in the breadcrumb trail to see a view with more information about that category.

This chapter contains information about the following views:

Symmetrix systems view (Root Cause Analysis)

The **Root Cause Analysis** view provides high-level data about all storage systems. The specific information displayed depends on the models of the selected storage systems. Some or all of the following information is displayed:

- **Symmetrix ID** — The Symmetrix ID for each supported system in the domain.
- **Alerts** — The number and severity of alerts associated with each Symmetrix system.
- **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 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.

System

Alerts view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about alerts. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

FE Directors view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about front-end directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

BE Directors (DA) view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about back-end directors (DA). The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

External Directors view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about external directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of each external director.
- **% Busy** — The percent of time that the director is busy.
- **IOs/sec** — The number of IOs (per second) handled by the director.
- **Reqs/sec** — The average number of requests performed by the director per second.
- **MBs Read/sec** — The average number of MBs read by the director per second.
- **MBs Written/sec** — The average number of MBs written by the director per second.

RDF Directors view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 is a generic SRDF term meaning a full track that comes from background data movement. These copies can be generated by a number of different things, including if the RDF device is in adaptive copy mode. Adaptive copy is included in this term, but not exclusively. For example, if SRDF is suspended, hosts continue to write to the R1 and then SRDF is resumed, the resulting R1->R2 copy of tracks (due to invalids) is a "copy."
- **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.

IM Directors view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about IM directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the IM director.
- **Page In Time** — .
- **Page Out Time** — .

EDS Directors view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about EDS directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of each EDS director.
- **Random Read Misses** — The number of random read IOs that were misses.
- **Random Read Misses Mbytes** —
- **Random Write Misses** — The number of random write IOs that were misses.
- **Random Write Misses Mbytes**—

Cache Partitions view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about cache partitions. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Boards view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about boards. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the board.
- **% Busy** — The percent of time the board is busy.
- **Outbound MBs/sec** —
- **Inbound MBs/sec** —
- **Max Speed MBs/sec** —

Disk Technologies view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about disk technologies. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The disk technology identifier.
- **Name** — The disk technology name.
- **% Busy** — The percent of time the disk technology group is busy.
- **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 technology group.
- **Avg Read Response time (ms)** — The average time it took the disk technology group to serve one read command.
- **Writes/sec** — The number of write operations performed by the disk technology group per second.
- **MBs Written/sec** — The number of write MBs performed by the disk technology group per second.
- **Avg Write Response Time (ms)** — The average time it took the disk technology group to serve one write command.
- **Total Capacity (GB)** — The total capacity of the disk technology group.
- **% Used Capacity**— The used capacity of the disk technology group.
- **Members** — The number of members in the disk technology group.
- **Model** — The vendor model number of the disk.

Events view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about events. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **Sequence Number** — .
- **Timestamp** — .
- **Function Class** — .
- **Action Code** — .
- **Activity ID** — .
- **Message** — .
- **Application ID** — .
- **Host ID** — .
- **Vendor ID** — .
- **User name** — .
- **Offset** — .

Storage

Storage Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Device Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about device groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 storage system to serve one read IO for this group.
- **Write Response Time (ms)** — The average time that it took the storage 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.

Databases view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about databases. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the database.
- **Host** —
- **Env Name** —
- **Host IOs/sec** — Host operations performed each second by the group.
- **Read Response Time (ms)** — The average time that it took the storage system to serve one read IO for this group.
- **Write Response Time (ms)** — The average time that it took the storage system to serve one write IO for this group.
- **% Hit** — Percentage of I/O operations 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.
- **% Writes** — Percentage of IO operations that were writes.

- **% Read Miss** — Percentage of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
- **Allocated Capacity (GB)** —
- **Total Capacity (GB)** —
- **Members** —

Thin Pools view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about virtual pools. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Disk Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about disk groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.
- **Writes/sec** — The number of write operations performed by the disk group per second.
- **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.

External Disk Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about external disk groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 storage 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 storage system to perform write operations.
- **Total Capacity (GB)** — The total available capacity in GBs for disks in the external disk group.
- **% Used Capacity** — The total percentage used capacity in GBs for disks in the external disk group.
- **Members** — The number of members in the external disk group.

SRPs view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about SRPs. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the SRP.
- **Name** — The name of the SRP.
- **Host IOs/sec** — The number of IOs per second transferred from the host..
- **Host MB/sec** — The number of host MBs written and read by all of the storage volumes each second.
- **BE Reqs/sec** — The number of read or write requests performed each second between the cache and the director.
- **BE MBs Transferred/sec** — The MBs transferred per second to the disk directors.
- **Avg Response Time (ms)** — The average time it takes the SRP to service IO.

Data Protection

RDFA Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about RDF/A groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — Displays the RDF/A group number.
- **Name** — Displays the RDF/A group name.
- **Host IOs/sec** — The total IOs per second for the RDF/A group.
- **Duration of Last Cycle** — The number of seconds for the last cycle in the session.
- **RDF R1 to R2 MBs/sec** — The throughput in MB 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.
- **% Compressed MBs Sent** — The percent of total MBs sent by the director that were compressed MBs.

RDFS Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about RDF/S groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The RDF/S group number.
- **Name** — The RDF/S group name.
- **Response Time (ms)** — The average time it took the storage system to serve one IO for the RDF/S group.
- **Host 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.

Snap Pools view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about snap pools. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 Read/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.

DSE Pools view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about DSE pools. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

FAST

FAST VP Policies view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about FAST policies. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Disk Group Tiers view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about disk group tiers. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 (ms)** — The average time it takes the storage 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 (ms)** — The average time it takes the storage 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.

Virtual Pool Tiers view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about virtual pool tiers. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name assigned to this storage tier.
- **BE Reqs/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.

Storage Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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

Hosts

Hosts view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about hosts. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the masking group.
- **Host IOs/sec** — The number of host operations performed each second by the host.
- **Host MBs/sec** — Number of MBs per second being processed for the host.
- **Read RT (ms)** — The average time it takes to serve one read IO.
- **Write RT (ms)** — The average time it takes to serve one write IO.

Initiators view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about initiators. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the masking group.
- **Host IOs/sec** — The number of host operations performed each second by the initiator.
- **Host MBs/sec** — Number of MBs per second being processed for the initiator.
- **Read RT (ms)** — The average time it takes to serve one read IO.
- **Write RT (ms)** — The average time it takes to serve one write IO.

Masking Views view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about masking views. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the masking group.
- **Host IOs/sec** — The number of host operations performed each second by the masking group.
- **Host MBs/sec** — Number of MBs per second being processed for the specific masking group.
- **Read Response-Time (ms)** — The average time it takes to serve one read IO.
- **Write Response-Time (ms)** — The average time it takes to serve one write IO.
- **Capacity (GB)** — The total capacity of the masking group, in GBs.

Port Groups view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about port groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the port group.
- **% Busy** — The percent of time the port group is busy.
- **Host IOs/sec** — The number of host IO operations performed each second, including writes and random and sequential reads.
- **Host MBs/sec** — The number of host MBs read each second.
- **Host IO Limit Delayed IOs/sec** — The sum of the IOs that were delayed due to the port group's IO limits and bandwidth limits.

Host IO Limit by SG view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about host IO limits set for the storage group. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the storage group.
- **Host IOs/sec** — Number of IOs the FE port is completing for the storage group.
- **Read Response-Time (ms)** — The average time it takes to serve one read IO for this storage group.
- **Write Response-Time (ms)** — The average time it takes to serve one write IO for this storage group.
- **IO Limit % Delayed IOs** — The host IO delayed IOs divided by the IO count.
- **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 host IO delayed IOs per second.

Host IO Limit by FE view (Root Cause Analysis)

The **Root Cause Analysis** view provides detailed data about host IO limits sorted by the front-end director. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the director.
- **Host IOs/sec** — Number of IOs the FE port is completing for the specific storage group.
- **Host IO Limit IOs/sec** — Number of IOs the FE port is completing for the specific storage group.
- **Host MBs/sec** — Number of MBs per second being processed for the specific storage group.
- **Host IO Limit MBs/sec** — The Host IO delayed IOs divided by the IO count.
- **Host IO Limit % Delayed IOs** — The Host IO delayed IOs divided by the IO count.

Trending & Planning views

You can use the Trending & Planning views to view historical data, showing long-term trends, to plan for future needs.

To open a Trending & Planning view:

1. Click **Performance › Analyze**.
2. Click **Trending & Planning**.
3. Click the **Symmetrix ID** of the storage system for which you want to view an analyze view.
4. Select a category from the drop-down menu in the breadcrumb trail to see a view with more information about that category.

This chapter contains information about the following views:

Symmetrix systems view (Trending & Planning)

The **Trending & Planning** view provides high-level data about all storage systems. The specific information displayed depends on the models of the selected storage systems. Some or all of the following information is displayed:

- **Symmetrix ID** — The Symmetrix ID for each supported system in the domain.
- **Alerts** — The number and severity of alerts associated with each storage system.
- **Host IOs/sec** — The number of host IO operations performed each second by all volumes, including writes and random and sequential reads.
- **Host MBs/sec** — The number of host MBs written and read by all of the 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 storage 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 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.

System

Alerts view (Trending & Planning)

The **Trending & Planning** view provides detailed data about alerts. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

FE Directors view (Trending & Planning)

The **Trending & Planning** view provides detailed data about front-end directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

BE Directors (DA) view (Trending & Planning)

The **Trending & Planning** view provides detailed data about back-end directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

External Directors view (Trending & Planning)

The **Trending & Planning** view provides detailed data about external directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of each external director.
- **% Busy** — The percent of time that the director is busy.
- **IOs/sec** — The number of IOs (per second) handled by the director.
- **Reqs/sec** — The average number of requests performed by the director per second.
- **MBs Read/sec** — The average number of MBs read by the director per second.
- **MBs Written/sec** — The average number of MBs written by the director per second.

RDF Directors view (Trending & Planning)

The **Trending & Planning** view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 is a generic SRDF term meaning a full track that comes from background data movement. These copies can be generated by a number of different things, including if the RDF device is in adaptive copy mode. Adaptive copy is included in this term, but not exclusively. For example, if SRDF is suspended, hosts continue to write to the R1 and then SRDF is resumed, the resulting R1->R2 copy of tracks (due to invalids) is a “copy.”
- **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.

IM Directors view (Trending & Planning)

The **Trending & Planning** view provides detailed data about IM directors. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the IM director.

- **Page In Time** — .
- **Page Out Time** — .

EDS Directors view (Trending & Planning)

The Trending & Planning view provides the following information about EDS directors:

- **ID** — The ID of each EDS director.
- **Random Read Misses** — The number of random read IOs that were misses.
- **Random Read Misses Mbytes** —
- **Random Write Misses** — The number of random write IOs that were misses.
- **Random Write Misses Mbytes**—

Cache Partitions view (Trending & Planning)

The **Trending & Planning** view provides detailed data about cache partitions. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Boards view (Trending & Planning)

The **Trending & Planning** view provides detailed data about boards. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the board.
- **% Busy** — The percent of time the board is busy.
- **Outbound MBs/sec** —
- **Inbound MBs/sec** —
- **Max Speed MBs/sec** —

Disk Technologies view (Trending & Planning)

The **Trending & Planning** view provides detailed data about disk technologies. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The disk technology identifier.
- **Name** — The disk technology name.

- **% Busy** — The percent of time the disk technology group is busy.
- **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 technology group.
- **Avg Read Response time (ms)** — The average time it took the disk technology group to serve one read command.
- **Writes/sec** — The number of write operations performed by the disk technology group per second.
- **MBs Written/sec** — The number of write MBs performed by the disk technology group per second.
- **Avg Write Response Time (ms)** — The average time it took the disk technology group to serve one write command.
- **Total Capacity (GB)** — The total capacity of the disk technology group.
- **% Used Capacity**— The used capacity of the disk technology group.
- **Members** — The number of members in the disk technology group.
- **Model** — The vendor model number of the disk.

Events view (Trending & Planning)

The **Trending & Planning** view provides detailed data about events. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **Sequence Number** — .
- **Timestamp** — .
- **Function Class** — .
- **Action Code** — .
- **Activity ID** — .
- **Message** — .
- **Application ID** — .
- **Host ID** — .
- **Vendor ID** — .
- **User name** — .
- **Offset** — .

Storage

Storage Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Device Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about device groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 storage system to serve one read IO for this group.
- **Write Response Time (ms)** — The average time that it took the storage 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.

Databases view (Trending & Planning)

The **Trending & Planning** view provides detailed data about databases. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the database.
- **Host** —
- **Env Name** —
- **Host IOs/sec** — Host operations performed each second by the group.
- **Read Response Time (ms)** — The average time that it took the storage system to serve one read IO for this group.
- **Write Response Time (ms)** — The average time that it took the storage system to serve one write IO for this group.
- **% Hit** — Percentage of I/O operations 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.
- **% Writes** — Percentage of IO operations that were writes.
- **% Read Miss** — Percentage of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
- **Allocated Capacity (GB)** —
- **Total Capacity (GB)** —
- **Members** —

Thin Pools view (Trending & Planning)

The **Trending & Planning** view provides detailed data about virtual pools. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Disk Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about disk groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.
- **Writes/sec** — The number of write operations performed by the disk group per second.
- **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.

External Disk Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about external disk groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 storage 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 storage system to perform write operations.

- **Total Capacity (GB)** — The total available capacity in GBs for disks in the external disk group.
- **% Used Capacity** — The total percentage used capacity in GBs for disks in the external disk group.
- **Members** — The number of members in the external disk group.

SRPs view (Trending & Planning)

The **Trending & Planning** view provides detailed data about SRPs. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the SRP.
- **Name** — The name of the SRP.
- **Host IOs/sec** — The number of IOs per second transferred from the host..
- **Host MB/sec** — The number of host MBs written and read by all of the storage volumes each second.
- **BE Reqs/sec** — The number of read or write requests performed each second between the cache and the director.
- **BE MBs Transferred/sec** — The MBs transferred per second to the disk directors.
- **Avg Response Time (ms)** — The average time it takes the SRP to service IO.

Data Protection

RDFA Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about RDF/A groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — Displays the RDF/A group number.
- **Name** — Displays the RDF/A group name.
- **Host IOs/sec** — The total IOs per second for the RDF/A group.
- **Duration of Last Cycle** — The number of seconds for the last cycle in the session.
- **RDF R1 to R2 MBs/sec** — The throughput in MB 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.

- **% Compressed MBs Sent** — The percent of total MBs sent by the director that were compressed MBs.

RDFS Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about RDF/S groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The RDF/S group number.
- **Name** — The RDF/S group name.
- **Response Time (ms)** — The average time it took the storage system to serve one IO for the RDF/S group.
- **Host 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.

Snap Pools view (Trending & Planning)

The **Trending & Planning** view provides detailed data about snap pools. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 Read/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.

DSE Pools view (Trending & Planning)

The **Trending & Planning** view provides detailed data about DSE pools. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

FAST

FAST VP Policies view (Trending & Planning)

The **Trending & Planning** view provides detailed data about FAST policies. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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.

Disk Group Tiers view (Trending & Planning)

The **Trending & Planning** view provides detailed data about disk group tiers. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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 (ms)** — The average time it takes the storage 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 (ms)** — The average time it takes the storage 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.

Virtual Pool Tiers view (Trending & Planning)

The **Trending & Planning** view provides detailed data about virtual pool tiers. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name assigned to this storage tier.
- **BE Reqs/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.

Storage Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **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

Hosts

Hosts view (Trending & Planning)

The **Trending & Planning** view provides detailed data about hosts. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the masking group.
- **Host IOs/sec** — The number of host operations performed each second by the host.
- **Host MBs/sec** — Number of MBs per second being processed for the host.
- **Read RT (ms)** — The average time it takes to serve one read IO.
- **Write RT (ms)** — The average time it takes to serve one write IO.

Masking Views view (Trending & Planning)

The **Trending & Planning** view provides detailed data about masking views. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the masking group.
- **Host IOs/sec** — The number of host operations performed each second by the masking group.
- **Host MBs/sec** — Number of MBs per second being processed for the specific masking group.
- **Read Response-Time (ms)** — The average time it takes to serve one read IO.
- **Write Response-Time (ms)** — The average time it takes to serve one write IO.
- **Capacity (GB)** — The total capacity of the masking group, in GBs.

Port Groups view (Trending & Planning)

The **Trending & Planning** view provides detailed data about port groups. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The ID of the port group.
- **% Busy** — The percent of time the port group is busy.

- **Host IOs/sec** — The number of host IO operations performed each second, including writes and random and sequential reads.
- **Host MBs/sec** — The number of host MBs read each second.
- **Host IO Limit Delayed IOs/sec** — The sum of the IOs that were delayed due to the port group's IO limits and bandwidth limits.

Host IO Limit by SG view (Trending & Planning)

The **Trending & Planning** view provides detailed data about host IO limits set for the storage group. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the storage group.
- **Host IOs/sec** — Number of IOs the FE port is completing for the storage group.
- **Read Response-Time (ms)** — The average time it takes to serve one read IO for this storage group.
- **Write Response-Time (ms)** — The average time it takes to serve one write IO for this storage group.
- **IO Limit % Delayed IOs** — The host IO delayed IOs divided by the IO count.
- **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 host IO delayed IOs per second.

Host IO Limit by FE view (Trending & Planning)

The **Trending & Planning** view provides detailed data about host IO limits sorted by the front-end director. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the director.
- **Host IOs/sec** — Number of IOs the FE port is completing for the specific storage group.
- **Host IO Limit IOs/sec** — Number of IOs the FE port is completing for the specific storage group.
- **Host MBs/sec** — Number of MBs per second being processed for the specific storage group.
- **Host IO Limit MBs/sec** — The Host IO delayed IOs divided by the IO count.
- **Host IO Limit % Delayed IOs** — The Host IO delayed IOs divided by the IO count.

Initiators view (Trending & Planning)

The **Trending & Planning** view provides detailed data about initiators. The specific information displayed depends on the model of the selected storage system. Some or all of the following information is displayed:

- **ID** — The name of the masking group.
- **Host IOs/sec** — The number of host operations performed each second by the initiator.
- **Host MBs/sec** — Number of MBs per second being processed for the initiator.
- **Read RT (ms)** — The average time it takes to serve one read IO.
- **Write RT (ms)** — The average time it takes to serve one write IO.

Charts view

This chapter contains the following:

Using the UI elements in the Charts view

Depending on the current charts, some of the following icons are found in the Charts view. Each icon provides a description when you mouse over it.

UI element	Description
	You can use this menu to change the time range.
	Edits the currently selected filter criteria.
	You can use this menu to select the metric (in the KPI list) by which the list of top *n* most active instances is sorted.
	Draws a time series chart using the selected elements. Time series charts display the values charted over time.
	Draws a snapshot chart using the selected elements. A snapshot chart displays a single average value over time.
	Draws a threshold chart using the selected elements.

UI element	Description
	Adds the selected elements to the chart.
	Minimizes the chart.
	◆ **Chart Style** — Changes the chart style. Depending on the properties of the current chart, some or all of the following options are available: • **Line** • **Bar** • **Stacked Bar** • **Area** • **Pie** • **Table** ◆ **Data Format** — Displays the average or maximum metric value, or both, for the session. • **Average** • **Maximum** ◆ **Zoom** — Changes the zoom settings of the chart. The following options are available: • **Zoom In** • **Zoom Out** • **Reset Zoom** ◆ **Add to Chart** — Adds the selected metric to the current chart. ◆ **Trend Line** — Adds a trend line to the current chart. ◆ **Display Thresholds** — Adds a line to the chart to show the user-defined thresholds. If a threshold is not defined, this option is dimmed. ◆ **Save as Image** — Saves the current chart as a JPEG image.
	Maximizes the chart.
×	Removes the chart.
Save As	Lists the formats to which you can save the chart. Some or all of the following options are available: ◆ Dashboard

UI element	Description
	◆ Template ◆ JPG ◆ PNG ◆ To Clipboard ◆ PDF - Charts ◆ PDF - Chart & Data
Remove all Charts	Deletes all displaying charts.
1 column 2 column 3 column 4 column 1 column	Sets the number of columns for the chart view.
	Restores a maximized chart to the default chart size.
	Restores a minimized chart to the default chart size.
Forecast Points: 0	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.

Creating charts

The Charts 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.

When you select a particular system and/or component for which you want to create a chart, the instances per category are listed in order of how active they are, according to the default KPI displayed in the drop-down menu. The default KPI used depends on the component type but you can select a different KPI, as required.

To create charts:

1. Select **Performance › Monitor** Click **Charts** to open the **Charts** view.
2. Select an appropriate time range for the chart or specify a custom time range.

3. In the **Categories & Instances** section, select the systems or components for which you want to create a chart. Multi-select is allowed in the table.

   For some components (for example, **Disks** and **Volumes**), you are prompted to narrow the list of displayed components according to filter criteria (for example, disk technology). You can click **OK** without specifying filter criteria, or you can select filter criteria to apply and then click **OK**.

4. *Optional:* To sort instances by category by a different metric, select a KPI from the drop-down list of KPI metrics available for that category.

5. Select the metric(s) that you want to use from the **KPI Metrics** list or the **All** list. You can select multiple metrics, however, the more you select, the less readable the chart will be.

6. Click the icon of the chart type you want to create. The following options are available:

   - **Time series chart**
     This chart displays values charted over time.
   - **Snapshot chart**
     This chart displays a single average value over time.
   - **Threshold chart**
     This chart lists, in a table, any thresholds applicable to the selected metrics.

   Each chart has icons for properties (style, format, thresholds), minimization/maximization, delete, and refresh.

7. *Optional*: Customize the chart. For more information about how to do this, see Customizing charts.

8. *Optional*: Customize the Chart view in which the chart is displayed. For more information about how to do this, see Customizing the Charts view.

Customizing charts

After you create a chart there several ways to customize that chart:

- See Filtering the information in a chart on the facing page.
- See Changing the style of a chart on the facing page.
- See Changing the data format of a chart on the facing page.
- See Displaying the Trend Line on a chart on page 618.
- See Adding a metric to a chart on the facing page.
- See Changing the zoom level of a chart on the facing page.
- See Displaying the Trend Line on a chart on page 618.
- See Saving a chart as an image on page 618.

- See Saving a chart as a CSV file on the next page.

Filtering the information in a chart

If the **Categories & Instances** specified for a chart support filtering, you can apply filter criteria to the components for which information is displayed in a chart. Click to apply filters, or edit existing ones.

The **Categories & Instances** that support filtering include **Disks** and **Volumes**.

Changing the style of a chart

The type of chart styles available depend on the metrics selected for the current chart. Some or all of the following chart styles are available:

- Area
- Bar
- Line
- Pie
- Stacked Bar
- Table

To change the chart style, click **› Chart Style** and select the chart style you want to apply to the current chart.

Changing the data format of a chart

By default, the Charts view displays the average value for each metric 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.

Displaying a threshold on a chart

If thresholds have been set for a metric, you can display those threshold values on the chart. Click **› Display Thresholds**. Threshold values are not displayed in table charts.

Adding a metric to a chart

To add additional metrics to a chart, select the metric in the **KPI Metrics** or **All** list, and click .

Changing the zoom level of a chart

When viewing a time series chart, you can zoom in to see the chart in more detail (**› Zoom › Zoom In**) or zoom out to view a larger time interval **› Zoom › Zoom Out**.

To view a particular interval in more detail, select the time interval on the graph and the zoom level of the chart automatically adjusts to display only that time interval.

To reset the zoom level of a chart, click › **Zoom** › **Reset Zoom**.

Displaying the Trend Line on a chart

To display the Trend Line on a chart, click › **Trend Line** and from the **Forecast Points** menu, select an appropriate number of data points to use for the trend line.

Saving a chart as an image

To save a chart as an JPEG image click › **Save as** . Select **JPG** or **PNG**. Then specify a file name for the image file.

Saving a chart as a CSV file

To save a chart as an CSV file, click › **Save as** › **CSV**. Then specify a file name for the CSV file.

Customizing the Charts view

There several ways to customize the way that charts are displayed on the Charts view:

- See "Saving the Charts view as a dashboard" on page 619
- See "Saving the Charts view as a template" on page 619
- Opening an image of the Charts view in a new browser tab
- See "Removing all of the charts from the Charts view" on page 619
- See "Adjusting the number of columns displayed in the Chart view" on page 619

Saving the Charts view as a dashboard

After you have customized the Charts view, you can save it as a dashboard. To do this. click **Save As › Dashboard**.

Saving the Charts view as a template

After you have customized the Charts view, you can save it as a template dashboard. To do this, click **Save As › Template**. For more information, see See "Creating a template dashboard from an Analyze view" on page 579

Opening an image of the Charts view in a new browser tab

After you have customized the Charts view, you can open a JPG image of the current chart in a new browser tab. From there, it can copied to the clipboard and pasted to another application, for example, into a presentation. To do this, click **Save As › To Clipboard**

Note This functionality is available only when using Firefox or Chrome web browsers.

Removing all of the charts from the Charts view

To remove all of the charts from the Charts view, click **Remove all Charts**.

Adjusting the number of columns displayed in the Chart view

To change the number of columns in which the charts in the Charts view are displayed, click the column icon drop-down menu.



Plan view

This chapter contains the following:

Using projection dashboards

The performance projection dashboard displays predicted future data, based on linear projection. You can use the information contained in the projection dashboard to assist when planning and to avoid potential problems.

The following projection dashboards are available:

- Array projection dashboard
- Thin pool projection dashboardThin pool projection dashboard

Select the storage system and category for which you want to view a projection dashboard. By default when initially opened, the projection dashboard displays a heatmap with the Array category of the first storage system.

You need a minimum of 10 days of data to produce a projection dashboard. If insufficient data is available, a message stating the number of remaining days before the projection dashboard can be produced is displayed.

- Save as — A drop-down menu containing some or all of the following options:
  - JPG — Saving dashboards and charts on page 562
  - PNG — Saving dashboards and charts on page 562
  - PDF Charts — Saving dashboards and charts on page 562
  - PDF Chart & Data — Saving dashboards and charts on page 562
- **Schedule Report** — Scheduling a report from the dashboard on page 572

Array projection dashboard

The array protection dashboard displays utilization data based on projections calculated over the selected time range.

You can configure the time ranges for the dashboard by selecting an appropriate time range from the drop-down menu. The following time ranges are available:

- Next week
- Next Month
- Next 3 Months
- Next 6 Months

By default, when initially opened, the Utilization Projection dashboard displays heatmap data for the next month.

Click on an element in the array to display a chart for the selected component showing the current utilization values and the projection for the selected time frame.

Thin pool projection dashboard

The Thin pools protection dashboard displays the following charts for the Thin pools in the array:

- All Thin Pools - Capacity Status
- Thin pool projection dashboard

Note Changing the time range for the dashboard changes the time range for the **% Used Capacity Projections** chart only. The **All Thin Pools - Capacity Status** chart always displays the current data.

All Thin Pools - Capacity Status

This chart displays a table containing current information for all Thin pools in the array. The table contains the following columns:

- **ID**—The ID of the Thin pool.

  You can click the ID of the Thin pool to navigate to the **Single Thin Pool EMC** dashboard in the **Monitor** view.

- **Days to Full**—The number of days until the capacity of the pool will be fully utilized.

  The maximum calculated value is 180 days. One of the following status icons is displayed according to the **Days to Full** value:

  - ‹ 7 days utilization remaining
  - ‹ 30 days utilization remaining
  - ‹ 90 days utilization remaining
  - › 90 days utilization remaining

- **Capacity Trend/Month**—The capacity growth rate of the Thin pool per month. Arrow icons indicate if the going up, down, or remaining stable.
- **Capacity**—The total capacity of the Thin pool and the percentage currently being used.

By default, the table is sorted by the **Days to Full** column.

% Used Capacity Projections

This chart displays the current and projected % used capacities for the currently selected Thin pool. When the chart is initially opened, the first Thin pool is displayed, by default.

You can configure the time ranges for this chart by selecting an appropriate time range using the slider control. The following time ranges are available:

- Next week
- Next Month
- Next 3 Months
- Next 6 Months

By default, when initially opened, the **% Used Capacity Projections** chart displays data for the next week.

Changing the time range for the dashboard changes the time range for this chart only. The **% Used Capacity Projections** chart always displays current data.

Reports view

This chapter contains the following:

Viewing performance reports

1. Select **Performance › 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 storage 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 Creating performance reports below.
- **Copy** — See Copying performance reports on page 625.
- **Edit** — See Editing performance reports on page 626.
- **Schedule** — See Scheduling performance reports on page 627.
- **Cancel Schedule** — See Cancelling a scheduled report on page 627.
- **Run Now** — See Running performance reports on page 627.
- **Delete** — See Deleting performance reports on page 628.
- **Navigate to Notifications** — See Configuring alert notifications on page 48.

When you click **Run Now**, you are prompted to choose the location for the generated reports.

Scheduled `XML`, `CSV`, and `BTP` reports are saved to this location:

install_dir`\SMAS\report\SPA\querydata`

Scheduled **UPV** reports are saved to this location:

`SMAS\backup\SPA`

Creating performance reports

To create a performance report:

1. Select **Performance › Reports** to open the **Reports** list view.

2. Click **Create** to open the **Create Report** dialog box.
3. In the General tab, enter a report **Name**.
4. *Optional*: Enter a report **Description**.
5. Select the **Symmetrix ID**.
6. Select a **Format** (XML, CSV, BTP, UPV).
   BTP reports are not available for VMAX3 systems.
7. Select the **Time Range**.

   If you select **Custom**, the Custom Time Selection dialog box displays for custom start and end times.
   The **Last Interval** option is available only for the **Root Cause Analysis** type. This option executes a report on the array and director categories at each Root Cause Analysis interval. In addition, every hour the system runs a cleanup job to remove Last Interval reports that are older than one hour.

   The BTP format is not available for **Last Interval** reports.
8. *Optional*: For UPV reports, specify whether to include RT traces.
9. From the **Generated Time Zone** menu, select the time zone to apply to times in the report.
10. For XML or CSV reports, add one or more queries to the report. Each of these reports must have at least one query.

    For more information about creating queries see Creating queries using the Create Query wizard below
11. *Optional*: In the Schedule tab, select the **Schedule** check box and configure values for each of the following options:
    - **First Run Time**
    - **Day(s) to run**
    - **Retention (Days)**
12. *Optional*: In the **Email** tab, specify the report email options.
13. Click **OK**.

Creating queries using the Create Query wizard

Before you begin

The **Create Query** wizard is launched from the Create Report dialog box. For more information about creating a report, see Creating performance reports on the previous page.

For some categories, you can configure the query to return the *n* most active instances. This feature is available for the following categories:

Category	Metrics
Storage Group	Host IOs/sec (default) Host MBs/sec Response Time

Storage Group Volumes	Host IOs/sec (default) Host MBs/sec Response Time
RFDS Group	Host IOs/sec (default) Host MB/Sec Response Time
RFDA Group	RDF R1 to R2 IOs/sec (default) RDF R1 to R2 MBs/sec RDF Response Time Host Writes/sec Host MBs Written/sec

To create a query using the Create Query wizard

1. When creating a report, in the Create Report dialog box, click **Add** to open the **Create Query** wizard. The **Query Name & Description** window is displayed.
2. Enter a query **Name**.
3. *Optional*: Enter a query **Description**.
4. Click **Next**. The **Category & Instance** page is displayed.
5. Select a **Category**. Depending on the category type you select, different options are presented.
6. Specify appropriate **Instances**. Some or all of the following options are available:
   - **All**—Displays all instances.
   - **Specific instances**—Displays the specific instances you select. Select them in the **Available** list and click **›** to move them to the **Selected** list. You can use multi-select for multiple instances.
   - **Most Active**—Displays the top *n* most active instances. You can specify a value for n and the metric by which you want to sort the displayed instances.
7. Click **Next**. The **Metrics** window is displayed.
8. Select a **Data Format**. The options are **Average** and **Maximum**.
9. Select the **Metrics** (for the category and data format you selected on the **Category & Instance** page of the wizard). You can use multi-select for multiple metrics.
10. Click **Finish**.

Copying performance reports

Before you begin:

You require an existing report in the **Reports** list view.

To copy a performance report:

1. Select **Performance › Reports** to open the Reports list view.
2. Select a report and click **Copy** to open the Copy Report dialog box.
3. Update the report settings, as appropriate. You can edit the following data:
   - **Name**
   - **Description**
   - **Symmetrix ID**
   - **Format**
   - **Time Range**
   - **Generated Time Zone**
4. In the **Queries** section, do one of the following:
   a. Click **Add** to open the **Create Query** wizard and create a 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 A report must have at least one query.

5. *Optional*: In the Schedule tab, configure the schedule settings for the report.
6. *Optional*: In the Email tab, configure the email settings for the report.
7. Click **OK**.

   For more information about reports, see Creating performance reports on page 623.

Editing performance reports

Before you begin:

You require an existing report in the Reports list view.

To edit a performance report:

1. Select **Performance › 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**.

   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 Creating performance reports on page 623.

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 › Reports** to open the **Reports** list view.
2. Select one or more reports and click **Schedule** to open the **Schedule Report** dialog box.
3. Select the **Schedule** check box.
4. Configure values for each of the following options:
   - **First Run Time**
   - **Day(s) to run**
   - **Retention (Days)**
5. Click **OK**.

Note Scheduled `XML`, `CSV`, and `BTP` reports are saved to this location:

```
SMAS\report\SPA\querydata
```

Scheduled **`UPV`** reports are saved to this location:

```
SMAS\backup\SPA
```

Cancelling 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 › Reports** to open the **Reports** list.
2. Select one or more reports and click **Cancel Schedule**.
3. Click **OK** in the confirmation message.

Running performance reports

Before you begin:

You require a configured performance report in the **Reports** list.

To run a performance report:

1. Select **Performance › Reports** to open the **Reports** list.
2. Select one or more reports and click **Run Now**.
3. Click **OK** on the **Save Report Results** confirmation.
4. Enter a location in the **Save As** dialog box, enter a file name and extension (if necessary), and click **Save**.

   The report is saved to the selected location.

Deleting performance reports

Before you begin:

You require an existing report in the **Reports** list view.

To view performance reports:

1. Select **Performance › Reports** to open the Reports list view.
2. Select the report you want to delete.
3. Click **Delete**.
4. Click **OK** on the confirmation message.

Settings

This chapter contains the following:

System Registrations

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 storage system performance registration and data collection policies.

   The following properties display:

   - **Symmetrix ID** — The ID of the registered storage system.
   - **Real Time** — An indicator that shows if the storage system is registered for Real Time data.
   - **Root Cause Analysis** — An indicator that shows if the storage system is registered for Root Cause Analysis data.
   - **Collection Recovery (Hours)** — The amount of data that the storstpd daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
   - **Root Cause Analysis Interval (Minutes)** — The frequency of polling in the Root Cause Analysis view. The default 5 minutes is recommended.
   - **Server IP** — The IP address of the Unisphere for VMAX server.

   The following controls are available:

   - **View Details** — See Viewing system registration details below.
   - **Register** — See Registering storage systems on the next page.

Viewing system registration details

The system registration details view shows the default values for storage 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 storage system.
- **Real Time** — An indicator that shows if the storage system is registered for Real Time data.
- **Root Cause Analysis** — An indicator that shows if the storage system is registered for Root Cause Analysis data.
- **Collection Recovery (Hours)** — The amount of data that the storstpd daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
- **Root Cause Analysis Interval (Minutes)** — The frequency of polling in the Root Cause Analysis 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.

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**.

Registering storage systems

The Performance component of Unisphere for VMAX collects performance data only after you register your storage systems.

Only local storage systems can be registered for performance data collection.

When you register a storage system for performance data collection it takes at least two intervals (by default, 5 minutes each) before performance data begins to populate in the GUI and charts. Before this, the GUI and charts might display "NA" or remain blank.

To register a storage system:

1. Select the storage system.
2. Click **Performance › Settings › System Registrations** to open the **Systems Registrations** view.
3. Click **Register** to open the **Register** dialog box. The Symmetrix ID automatically populates the **Selected Symmetrix(s)** list.
4. Select the type of performance data to collect: **Root Cause Analysis** only, or both **Real Time** and **Root Cause Analysis**.

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 storstpd daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
   - **Root Cause Analysis Interval (Minutes)** — The frequency of polling in the Root Cause Analysis view. The default 5 minutes is recommended.
5. Make your changes and click **Apply**.

Performance 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 the Root Cause Analysis view.

To create a threshold:

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Click **Create** to open the **Create Performance Thresholds and Alert** dialog box.
3. Select the **Symmetrix ID**, **Category**, **Instance**, and **Metric**.

   For some categories, when selected, the **Instance** field defaults to **All**.

   If there is a default value for the **Metric**, it automatically displays in the **First Threshold** and **Second Threshold** fields.
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**.
2. Select the severity for the first threshold, the second threshold, or both. The available 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, an alert is initiated.

4. *Optional:* Select any additional configuration options. For some group categories, you can choose to enable for the alert for the individual components of the group, for example, when the **Disk Group** category is selected, you have the option to enable the alert for the disk.
5. 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 that the value has changed from the default.

To edit performance thresholds and alerts

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Navigate to the threshold or alert to be edited by selecting the appropriate category in the **Category** section and expanding the appropriate metric in the **Metric** column.
3. Select the item from the table and click **Edit** to open the **Edit Performance Threshold and Alert** dialog box.
4. Edit the threshold and alert settings, as appropriate.

   For more information about threshold and alert settings, see See "Performance Thresholds and Alerts" on page 631
5. Click **OK**.

Deleting performance thresholds and alerts

Before you begin:

You can delete only custom values. You cannot delete default thresholds.

To delete a performance threshold and alert:

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Navigate to the threshold or alert to be edited by selecting the appropriate category in the **Category** section and expanding the appropriate metric in the **Metric** column.
3. Select one or more rows and click **Delete**.

   A confirmation message displays.
4. In the **Delete Threshold and Alert** confirmation dialog box, click **Delete**.

Viewing thresholds and alerts

To view thresholds and alerts:

1. Select **Performance › Settings › Performance Thresholds and Alerts** to open the **Performance Thresholds and Alerts** list.
2. Select the storage system.

3. In the **Category** section, select the category for which you want to view the configured thresholds and alerts.

   The thresholds and alerts configured for that category are displayed, according to metric. By default, the KPI tab displays. To view all metrics, click **All**.

   Any metrics that include a custom threshold or alert are highlighted with a tick mark in the **Custom** column.

The some of the following properties display:

- **Metric** — The storage system identifier.
- **KPI** — Indicates (with a tick mark) if the metric is a KPI.
- **Custom** — Indicates (with a tick mark) whether a custom threshold or alert has been set.
- **First Threshold** — The user-defined first threshold.
- **Second Threshold** — The user-defined second threshold.

The following controls are available:

- **Create** — See Performance Thresholds and Alerts on page 631.
- **Edit** — See Editing a performance threshold and alert on the previous page.
- **Delete** — See Deleting performance thresholds and alerts on the previous page.
- **Notifications** — See Configuring alert notifications on page 48.

Configuring email notification

You can configure email addresses to which notifications, alerts, and reports are sent. You can configure a single email address for all notification instances, or you can use different email addresses for different notifications on different storage systems.

To set up email notification:

1. Click **Home › Administration › Alert Settings › Notifications**. The **Notifications** window is displayed.
2. In the **Email** section, click **Enable**. The **Configure Email** window is displayed.
3. In the **From** section specify the following details:
   - **Sender Email**
   - **IP Address/Host**
   - **Server Port**
4. In the **Recipient Email** field of the **To** section, specify the address you want to enable for all email notifications.
5. Click **OK**.

To add an email address and apply it to all notification instances:

1. Click **Home › Administration › Alert Settings › Notifications**. The **Notifications** window is displayed.

2. In the **Email** section, click **Enable**. The **Configure Email** window is displayed.
3. In the **Recipient Email** field of the **To** section, specify the address you want to enable for all email notifications.

Note

If you configure a new email address in this way, any previously configured addresses are overwritten.

4. Click **OK**.

To configure email addresses for individual notifications:

1. Click **Home › Administration › Alert Settings › Notifications**. The **Notifications** window is displayed.
2. Click the **Edit** icon for the storage system for which you want to edit the email notification settings. The three individual email fields become editable.
3. Configure the email addresses for individual notifications, as appropriate.
4. *Optional*: Configure the notification levels by clicking the appropriate icons in the **System Level** and **Performance Level** columns.
5. Click **OK**.

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 storage 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 storage system IDs display if you are in **All Symmetrix** mode. If you have already selected a single storage system, this option is read-only.
4. Select the **Date** and **Trace 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.
8. Click **OK**.

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 storage systems.

There must be a scheduled Real Time Trace showing in the **Real Time Traces** list. Only scheduled (not yet executed) traces can be modified.

To edit a trace:

1. Select **Performance › Settings › Real Time Traces** to open the the **Real Time Traces** list.
2. Select a trace from the list and click **Edit** to open the **Edit Real Time Trace** dialog box.
3. Update the **Retention** and/or **Description** values. These are the only values that can be changed.
4. Click **OK**.

Deleting a Real Time trace

Before you begin:

You require 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.

To view a real time trace:

1. 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 list can be for a specific storage system, or all storage systems.

The following properties display:

- **Symmetrix ID** — The storage 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 for which the trace is be saved.
- **Description** — A user-defined description of the trace.

The following controls are available:

- **Create** — See Real Time Traces on page 634.
- **Edit** — See Editing a Real Time Trace on the previous page.
- **Delete** — See Deleting a Real Time trace on the previous page.

Databases

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 storage system identifier.
- **Status** — The status of the performance database.
- **Load Status** — The status of a database action (backup, restore).
- **Oldest Trending & Planning Data** — Date and time of the earliest Trending & Planning data.
- **Oldest Root Cause Analysis Data** — Date and time of the earliest Root Cause Analysis data.
- **Last Backup** — The last backup taken (in days).
- **Last Retention** — The last time data was saved (in hours).
- **Scheduled Backup** — Time of the scheduled backup.

The following controls are available:

- **Backup** — See Backing up a database below.
- **Restore** — See Restoring a database below.
- **View Details** — See Viewing database details on the next page.
- **Cancel Schedule** — See Canceling a scheduled database backup on page 639.
- **Delete Database** — See Deleting databases on page 639.
- **Remove Backup File** — See Removing database backup files on page 640.

Backing up a database

The backup database option is available for one or more storage systems, regardless of their registration status. By default, only Trending & Planning (Historical) data is backed up.

To back up the performance database:

1. Select **Performance › Settings › Databases** to open the **Databases** list.
2. Select a database from the list and click **Backup** to open the **Database Backup Wizard**.
3. Select the appropriate **Symmetrix ID**.
4. Enter a backup **File Name(s)**.
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 page 640..

Restoring a database

You can restore the performance database of a storage system regardless of whether it is registered.

Before you begin:

There must be an existing database file.

Restoring a performance database will result in the contents of the current performance database being overwritten. Ensure that you make a backup of the current database, if appropriate to do so.

Do not attempt to restore a database while another restore operation is in progress.

To restore a performance database:

1. Select **Performance › Settings › Databases** to open the **Databases** list.
2. Select the storage system.
3. From the **Select backup file** menu select a database to be restored.
4. Select **Restore Diagnostic data if available**if you want to restore Root Cause Analysis (Diagnostic) data.
5. Select **Restore Real Time traces if available**if you want to restore Real Time traces.
6. Click **OK**.

Viewing database details

Before you begin:

You require one or more storage 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 storage system identifier.
   - **Registered to Collection** — Shows the performance views available. Possible values are **Real Time** and **Root Cause Analysis**, 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 Trending & Planning data was derived from Root Cause Analysis data.
     - **Last retention time** — Date and time of last retention.
     - **Last real time retention time** — Date and time of last Real Time retention.
   - **Trending & Planning**
     - **Oldest available data** — Date and time of first available Trending & Planning data.
     - **Latest available data** — Date and time of latest Trending & Planning data.
     - **Set Trending & Planning data retention (months)** — Number of months to save Trending & Planning data. Possible values are 6 to 36 months. This option can be changed. After any change, click **Apply**.
   - **Root Causes Analysis**
     - **Oldest available data** — Date and time of first available Root Cause Analysis data.
     - **Latest available data** — Date and time of latest Root Cause Analysis data.

- **Set Root Cause Analysis 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 Root Cause Analysis** — Indicates whether the last day of Root Cause Analysis 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.

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** list.
2. Select a database and click **Cancel Schedule**.
3. Click **OK** on the confirmation message.

Deleting databases

The database delete action can only be performed on storage systems that are not registered. The delete action removes all references to the storage systems in the master database and removes the data. Backup files are not removed.

Before you begin:

You require one or more storage 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 the data you want to delete. The following options are available:
   - Delete database only
   - Delete database & backup files
   - Delete database & backup files and cancel scheduled tasks
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.
3. Select a database and click **Remove Backup Files**. The **Remove Backup File(s)**dialog box is displayed.
4. Select the file(s) to delete from the list of backup files and click **OK**.

Metrics

Viewing metrics

Use the **Metrics** list view to view and manage metrics.

To view metrics:

1. Select **Performance › Settings › Metrics** to open the **Metrics** list view.
2. Select a metric category from the **Category** menu.

   You can filter the list of displayed metrics. For more information about filtering, refer to Filtering performance data.

   The following properties display:

   - **Metric** — The ID of the performance metric.
   - **KPI** — An indicator that shows if the metric is a KPI.
   - **Hide** — An indicator that shows if the metric is hidden.

   The following controls are available:

   - **Save** — Editing metrics
   - **Restore Defaults** — Restores the default metric settings for the selected category.

Editing metrics

You can edit metrics to set, or unset, them as KPIs. Additionally, you can hide a KPI so that it does not appear on the UI. When a metric is not set as a KPI, it can be viewed under **All** in the Charts view.

To edit metrics:

1. Select **Performance › Settings › Metrics** to open the **Metrics** list view.
2. Select a metric category from the **Category** menu.
3. Do one of the following:
   - Edit the KPI status of one or more metrics:
     a. Select or clear the **KPI** checkbox.
     b. Click **Save**.
   - Edit the visibility of one or more metrics:
     a. Select or clear the **Hide** checkbox.
     b. Click **Save**.
   - Click **Restore Defaults** to restore default metric settings for all metrics in the selected category.

Metrics

This chapter contains the following:

Array metrics

The following table lists all metrics that are available for arrays.

Balance metrics measure hardware components balance. It is calculated as the difference between the maximum and minimum values for "% Busy" for the component instances. In a balanced system the balance metrics should be relatively low.

Table 8: Array metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second by all volumes, including writes and random and sequential reads.
Version WP Count	
Total Number & R1 Sessions	
Number of R1 Msc Active Sessions	
Number of R2 Cache Slots in Use	
Total Number of R2 Sessions	
Number of R1 Msc Sessions	
Number of R2 Msc Active Sessions	
Number of R1 Cache Slots in Use	
Number of R2 Msc Sessions	
RDFA WP Count	
Number of R2 Active Sessions	
Total Number of RDF Sessions	
Host Reads/sec	The number of host read operations performed each second by all

Metric	Description
	volumes.
Number of R1 Active Sessions	
Host Writes/sec	The number of host write operations performed each second by all volumes.
Host MBs/sec	The number of host MBs written and read by all of the volumes each second.
Host MBs Read/sec	The number of host MBs written and read by all of the volumes each second.
Host MBs Written/sec	The number of host MBs written by all of the 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.
Read RT (ms)	The average time it took the disk to serve one read command.
Write RT (ms)	The average time it took the disk to serve one write command.
% Reads	The percent of total read IO operations performed each second by all of the volumes.
% Writes	The percent of total write IO operations performed by all of the volumes.
% Hit	The percent of IO operations performed by all of the 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.
BE lOs/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

Metric	Description
	cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific volume was reached.
System 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.
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	The number of XtremSW Cache Dedup hits per second.

Metric	Description
Dedup Hits/sec	
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 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	
XtremSW Cache Read RT (ms)	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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 Read Misses/sec	
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	
FE Balance	

Metric	Description
Cache Balance	
SATA Balance	
FC Balance	
EFD Balance	
DX Balance	
Queue Depth Utilization	A weighted sum of the utilizations according to the number of IOs in each bucket. Queue depth range metrics are used in the calculation. Calculation: Sx=1..9 [(Avg Queue Depth Range X - 1) / Avg Queue Depth Range X] * (Queue Depth Count Range X / Queue Depth Count Total) * 100
DA Balance	
Allocated Capacity	
Critical Alert Count	
Warning Alert Count	
Information Alert Count	
Disk Utilization	
RDF Utilization	
BE Utilization	
FE Utilization	
Copy Slot Count	

BE Director (DA) metrics

The following table lists all metrics that are available for back-end directors.

Table 9: BE director (DA) metrics

Metric	Description
% Busy	The percent of time that a director is busy.
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.
% 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.
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.
Copy On First Access	Internal metric.
Clone Copy Read	
Clone Copy Write	
VLUN Migration Read	Internal metric.
VLUN Migration Write	Internal metric.
PHCO Rebuild Copy	Internal metric.

Metric	Description
Optimized Write	Internal metric.
PHCO Rebuild Read	Internal metric.
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 Reqs	The percentage of read requests that are compressed.
% Compressed Write Reqs	The percentage of write requests that are compressed.
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 MBs/sec	The total size of compressed MBs (read and write) per second.
Compressed Reqs/sec	The total number of compressed requests (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.
Syscall Time/sec	Total time spent processing all system calls
% 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 Logical Core 0	The percent of time that core 0 is idle.
% Idle Logical Core 1	The percent of time that core 1 is idle.
% 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.

BE Emulation metrics

The following table lists all metrics that are available for back-end emulation.

Table 10: BE emulation metrics

Metric	Description
Total Read Miss Time	
Total Write Time	A cumulative read time for calculating avg. read time
Total Prefetch Time	
Total Copy Read Time	
Total Copy Write Time	
Syscalls Time	
Management Time	
Rebuild Time	
IVTOC Time	
Software Compression Time	
Other Time	
CDI Time	
LP Task Time	
SYMMK Total Time	
SYMMK Idle Time	
Total Work Time	
% Busy	The percent of time the directory is busy.

BE Port metrics

The following table lists all metrics that are available for back-end ports.

Table 11: BE port metrics

Metric	Description
Reads/sec	
Writes/sec	
Host IOs/sec	
MBs Read/sec	
MBs Written/sec	
Host MBs/sec	
Avg IO Size (KB)	
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	

Board metrics

The following table lists all metrics that are available for boards.

The measurements reflect the transfer rate to and from BOSCO compared with the max bandwidth.

Utilization, as well as inbound and outbound speed, are measured across all processors, regardless of type.

Table 12: Board metrics

Metric	Description
Outbound MBs/sec	
Inbound MBs/sec	
Max Speed MBs/sec	
Utilization	Calculated value: ((transfer-rate-in + transfer-rate-out) / max-transfer-rate)

Cache Partition metrics

The following table lists all metrics 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 * % 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	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 Hit Ages 2	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 Hit Ages 3	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 Hit Ages 4	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.

Metric	Description
Cache Hit Ages 5	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 Hit Ages 6	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 Hit Ages 7	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 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	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Cache Slot Ages 2	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Cache Slot Ages 3	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Cache Slot Ages 4	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Cache Slot Ages 5	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Cache Slot Ages 6	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
Cache Slot Ages 7	A partition fall through under 8 seconds. A slot in this age range was removed or recycled as a new slot for other data.
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.
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.

Metric	Description
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).
Avg 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.

Core metrics

The following table lists all metrics that are available for cores.

Table 14: Core partition metrics

Metric	Description
% Busy	
SYMMK Idle Time	

DATA Volume metrics

The following table lists all metrics that are available for DATA volumes.

Table 15: DATA volume metrics

Metric	Description
BE Reqs/sec	Number of read/write requests each second performed by the disk directors to cache.
Parity Generation Reads	
Parity Generation Kbytes Read	
BE Read Reqs/sec	Number of read requests each second performed by the disk directors to cache.
BE Write Reqs/sec	Number of write requests each second performed by the disk directors to cache.
WP Count	The number of tracks currently in write pending mode for the volume.
BE MBs Transferred/sec	MBs read per sec + MBs written per sec
BE MBs Read/sec	Number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	Number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks Used/sec	Number of prefetched tracks used each second from the disk directors to the cache.
BE % Reads	
BE % Writes	
BE Prefetched MBs/sec	Total prefetched tracks each second from the disk directors to the cache.
BE Partial Sector Writes (KB)	
IO Density	
BE Optimize Writes (KB)	
BE XOR Reads	
BE XOR Reads (KB)	
BE Reads For Copy	
BE Reads For Copy (KB)	
BE Writes For Copy (KB)	

Metric	Description
Reads for Rebuild	
BE Reads For VLUN Migration	
BE Reads For VLUN Migration (KB)	
Kbytes Read for Rebuild	
BE Writes For VLUN Migration	
Writes for IVTOC	
Kbytes Written for IVTOC	
BE Writes For VLUN Migration (KB)	
BE Writes For Rebuild	
BE Writes For Rebuild (KB)	
Allocated Capacity (GB)	
Total Capacity (GB)	
Device Block Size	

Database metrics

The following table lists all metrics that are available for databases.

Table 16: Database metrics

Metric	Description
Host IOs/sec	Host operations performed each second by the group.
Host Reads/sec	Host read operations performed each second by the group.
RDF Reads/sec	
Host Writes/sec	Host write operations performed each second by the group.
RDF Writes/sec	
Host Hits/sec	Host read/write operations performed each second by the group that were immediately satisfied from cache.
Host Read Hits/sec	Host read operations performed each second by the group that were immediately satisfied from cache.
Host Write Hits/sec	Host write operations performed each second by the group that were immediately satisfied from cache.
Host Misses/sec	Host read/write operations performed each second by the group that could not be satisfied from cache.
Host Read Misses/sec	Host read operations performed each second by the group that were not satisfied from cache.
Host Write Misses/sec	Host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	Cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	Cumulative number of host MBs read per second by the group.
Host MBs Written/sec	Cumulative number of host MBs written per second by the group.
BE Reqs/sec	Number of read/write requests each second performed by the disk directors to cache.
BE Read Reqs/sec	Number of read requests each second performed by the disk directors to cache.
BE Write Reqs/sec	Number of write requests each second performed by the disk directors to cache.

Metric	Description
Read RT (ms)	
Write RT (ms)	
Read Miss RT (ms)	
Write Miss RT (ms)	
RDF/S Write RT (ms)	
% Reads	Percentage of IO operations that were reads.
% Writes	
% Read Hit	
% Write Hit	
% Read Miss	
% Write Miss	
WP Count	
Seq IOs/sec	
Seq Reads/sec	
Seq Writes/sec	
Seq Read Hits/sec	
Seq Read Miss/sec	
Seq Write Hits/sec	
Seq Write Misses/sec	
Random IOs/sec	
Random Reads/sec	
Random Writes/sec	
Random Read Hits/sec	
Random Write Hits/sec	
Random Read Misses/sec	
Random Write Misses/sec	
Avg IO Size (KB)	
Avg Read Size (KB)	
Avg Write Size (KB)	

Metric	Description
% Sequential IO	
% Seq Read	
% Seq Read Hit	
% Seq Read Miss	
% Seq Writes	
% Seq Write Hit	
% Seq Write Miss	
% Random IO	
% Random Read Hit	
% Random Write Miss	
% Random Read Miss	
% Random Write Hit	
Max WP Threshold	
BE MBs Transferred/sec	
BE MBs Read/sec	
BE MBs Written/sec	
BE Prefetched Tracks/sec	
BE Prefetched Tracks Used/sec	
BE Read Request Time (ms)	
BE Disk Read RT (ms)	
BE Read Task Time (ms)	
% Hit	
% Miss	
% Random Reads	
% Random Writes	
BE % Reads	
BE % Writes	
Write RT Count 5 to 8	
Write RT Count 8 to 14	

Metric	Description
Read RT Over 32	
Read RT Count 8 to 14	
Skew	
Read RT Count 0 to 1	
RDF Response Time (ms)	
Write RT Count 14 to 32	
Write RT Count Over 32	
Read RT 14 to 32	
Response Time (ms)	
Write RT 14 to 32	
Write RT Over 32	
Read RT Count 3 to 5	
Read RT Count 5 to 8	
Read RT Count 14 to 32	
Read RT Count Over 32	
Write RT Count 0 to 1	
Write RT Count 1 to 3	
Write RT Count 3 to 5	
Read RT Count 1 to 3	
Information Alert Count	
Critical Alert Count	
Warning Alert Count	
BE Partial Sector Writes (KB)	
IO Density	
BE Optimize Writes (KB)	
BE XOR Reads	
BE XOR Read (KB)	
BE Reads For Copy	
BE Reads For Copy (KB)	

Metric	Description
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 (MB)	
Total Capacity (GB)	
Device Block Size	
BE Prefetched MBs/sec	
XtremSW Cache Read Hits/sec	
RDF Read Hits/sec	
XtremSW Cache Write Hits/sec	
XtremSW Cache Reads/sec	
XtremSW Cache Writes/sec	
XtremSW Cache lOs/sec	
XtremSW Cache Skipped lOs/sec	
XtremSW Cache Dedup Hits/sec	
XtremSW Cache Dedup Reads/sec	
XtremSW Cache Dedup Writes/sec	
% XtremSW Cache Reads	
% XtremSW Cache Writes	
% XtremSW Cache Read Hits	
XtremSW Cache MBs Read/sec	
XtremSW Cache MBs Written/sec	
XtremSW Cache MBs/sec	

Metric	Description
XtremSW Cache Read RT (ms)	
XtremSW Cache Write RT (ms)	
XtremSW Cache RT (ms)	
XtremSW Cache Avg Read Size (KB)	
XtremSW Cache Avg Write Size (KB)	
XtremSW Cache Avg IO Size (KB)	
Optimized Read Misses/sec	
Optimized MBs Read Miss/sec	
RDF MBs Read/sec	
RDF Rewrites/sec	
RDF MBs Written/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	

Database by Pool metrics

The following table lists all metrics that are available for databases by pool.

Table 17: Database by pool metrics

Metric	Description
BE Reqs/sec	
Allocated Capacity	
IO Density	
BE Read Reqs/sec	
Destage Writes/sec	
BE MBs Transferred/sec	
BE MBs Read/sec	
Destage Write MBs/sec	
Prefetched Tracks/sec	
Pretetched MBs/sec	
Total Tracks	
Compressed Tracks	
% Compressed Tracks	

Device Group metrics

The following table lists all metrics that are available for device groups.

Table 18: Device group metrics

Metric	Description
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.
RDF Reads/sec	
RDF Writes/sec	
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.

Metric	Description
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.
Read RT (ms)	The average time that it took the Symmetrix to serve one read IO for this group.
Write RT (ms)	The average time that it took the Symmetrix to serve one write IO for this group.
Read Miss RT (ms)	The average time that it took the Symmetrix to serve one read miss IO for this group.
Write Miss RT (ms)	The average time that it took the Symmetrix to serve one write miss IO for this group.
RDF/S Write RT (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

Metric	Description
	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 Writes/sec	Write IO commands from a host not identified as part of a sequential stream.
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	
% Seq Read	Calculated value: 100 * (seq reads per sec/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*(seq writes per sec/total IOs per sec)
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.

Metric	Description
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Write Miss	Calculated value: 100 * (random write misses per sec / total IOs per sec)
% Random Read Hit	Calculated value: 100 * (random read hits per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec / total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
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	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 RT (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

Metric	Description
	requests.
BE % Writes	The percent of the back-end IO that were write requests.
RDF Response Time (ms)	
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 Writes (KB)	Internal metric
BE XOR Reads	Internal metric
BE XOR Read (KB)	Internal metric
BE Reads For Copy	Internal metric
BE Reads For Copy (KB)	Internal metric
BE Writes For Copy	Internal metric
BE Writes For Copy (KB)	Internal metric
BE Reads For VLUN Migration	Internal metric
BE Reads For VLUN Migration (KB)	Internal metric
BE Writes For VLUN Migration	Internal metric
BE Writes For VLUN Migration (KB)	Internal metric
BE Writes For Rebuild	Internal metric
BE Writes For Rebuild (KB)	Internal metric
BE RDF Copy	Internal metric
BE RDF Copy (MB)	Internal metric
Total Capacity (GB)	
Allocated Capacity (GB)	
Device Block Size	
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.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by

Metric	Description
	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.
XtremSW Cache Read RT (ms)	The average time it took the XtremSW Cache to serve one read for this volume.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write for this volume.
XtremSW Cache RT (ms)	The average time it took the XtremSW Cache to serve one IO for this volume.
XtremSW Cache Avg Read Size (KB)	The average size of a read served by the XtremSW Cache.
XtremSW Cache Avg Write Size (KB)	The average size of a write served by the XtremSW

Metric	Description
	Cache.
XtremSW Cache Avg IO Size (KB)	The average size of an IO (read or write) served by the XtremSW Cache.
Optimized Read Misses/sec	
RDF Rewrites/sec	
RDF MBs Read/sec	
Optimized MBs Read Miss/sec	Number of host MBs read each second directly from disks bypassing the cache.
RDF MBs Written/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	

Disk metrics

The following table lists all metrics that are available for disks.

Table 19: Disk 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.
IOs/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.
Response Time (ms)	The average response time for the reads and writes.
Read RT (ms)	The average time it took the disk to serve one read command.
Write RT (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.

Metric	Description
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 (GB)	The total capacity of the disk (GBs).
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 Bucket metrics

The following table lists all metrics that are available for disk buckets.

Table 20: Disk bucket metrics

Metric	Description
% Busy	The percent busy of 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 disk MBs read and written to disks in this bucket.
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.
Response Time (ms)	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.
Read RT (ms)	The average read response time.
Write RT (ms)	The average write response time.

Disk Group metrics

The following table lists all metrics 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.
Response Time (ms)	The average time it took the disk group to service IOs.
Read RT (ms)	The average time it took the disk group to serve one read command.
Write RT (ms)	The average time it took the disk group to serve one write command.
Seek Distance/sec	The number of hypervolumes that the disk head crossed (during all seeks) each second.
Total Capacity (GB)	The total capacity (in GB) of all the disks in the disk group.
Used Capacity (GB)	The total capacity (in GB) 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 +

Metric	Description
	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.)

Disk Group tier metrics

The following table lists all metrics that are available for disk group tiers.

Table 22: Disk group tier 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.
Response Time (ms)	The average time it took the disk group to service IOs.
Read RT (ms)	The average time it took the disk group to serve one read command.
Write RT (ms)	The average time it took the disk group to serve one write command.
Seek Distance/sec	The number of hypervolumes that the disk head crossed (during all seeks) each second.
Total Capacity (GB)	The total capacity (in GB) of all the disks in the disk group.
Used Capacity (GB)	The total capacity (in GB) 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 +

Metric	Description
	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.)

Disk Technology metrics

The following table lists all metrics that are available for disk technologies.

Table 23: Disk technology 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.
Response Time (ms)	The average time it took the disk group to service IOs.
Read RT (ms)	The average time it took the disk group to serve one read command.
Write RT (ms)	The average time it took the disk group to serve one write command.
Seek Distance/sec	The number of hypervolumes that the disk head crossed (during all seeks) each second.
Total Capacity (GB)	The total capacity (in GB) of all the disks in the disk group.
Used Capacity (GB)	The total capacity (in GB) 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 +

Metric	Description
	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.)

DSE Pool metrics

The following table lists all metrics that are available for DSE pools.

Table 24: 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 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 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 RT (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 (GB)	The enabled pool capacity in GBs.
Used Pool Capacity (GB)	The used pool capacity in GBs.
Allocated Pool Capacity (GB)	The allocated pool capacity in GBs.

DX Emulation metrics

The following table lists all metrics that are available for DX emulations.

Table 25: DX emulation metrics

Metric	Description
Total Read Miss Time	
Total Write Time	
Total Prefetch Time	
Total Copy Read Time	
Total Copy Write Time	
Syscalls Time	
Management Time	
Rebuild Time	
IVTOC Time	
Software Compression Time	
Other Time	
CDI Time	
LP Task Time	
SYMMK Total Time	
SYMMK Idle Time	
Total Work Time	
% Busy	

DX Port metrics

The following table lists all metrics that are available for DX ports.

Table 26: DX port metrics

Metric	Description
Reads/sec	
Writes/sec	
Host IOs/sec	
MSs Read/sec	
MSs Written/sec	
Host MBs/sec	
Avg IO Size (KB)	
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	

EDS Director metrics

The following table lists all metrics that are available for EDS directors.

Table 27: EDS director metrics

Metric	Description
Random Read Misses	
Random Read Misses Mbytes	
Random Write Misses	
Random Write Misses Mbytes	
% Busy	

EDS Emulation metrics

The following table lists all metrics that are available for EDS emulations.

Table 28: EDS director metrics

Metric	Description
Emulation Time	
IO Machine Time	
LP IO Machine Time	
Housekeeping Time	
LREP Control Time	
VP Time	
Copy Scanner Time	
Pyramid Time	
Scan Time	
Snow Spillover Time	
Remote Replication Time	
Syscall Time	
Watchdog Time	
Allocation Split Time	
FSM Time	
SYMMK Total Time	
SYMMK Idle Time	
Total Work Time	
% Busy	

External Director metrics

The following table lists all metrics that are available for external directors.

Table 29: External director metrics

Metric	Description
% Busy	The percent of time the directory is busy.
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.
Read RT (ms)	The calculated average response time for reads.
Write RT (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.)
Miss Reqs/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.

Metric	Description
% 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.
% 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.
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 0-9	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 0-9	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 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.
% Busy Logical Core 0	
Syscall Time/sec	
Avg RDF Write Time (ms)	
Avg WP Disconnect Time (ms)	

Metric	Description
Avg Read Miss Time (ms)	
Random MBs Written to non WP Slots/sec	
Random Writes to non WP Slots/sec	
Random MBs Written to WP Slots/sec	
Random Writes to WP Slots/sec	
% Busy Logical Core 1	
% Write Req Hit	The percent of write requests that were satisfied from cache.
Read RT Count Range 0 to Read RT Count Range Over 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 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.
Read RT 0 to 1 to Read RT Over 64	These buckets show the distribution of the total average read response time.
Write RT 0 to 1 to Write RT Over 64	These buckets show the distribution of the total average write response time.
Avg Write RT 16 to 32	
Write RT 32 to 64	
Write RT Over 64	
Host IO Limit IOs/sec	
Host IO Limit MBs/sec	
Queue Depth Utilization	

External Disk metrics

The following table lists all metrics that are available for external disks.

Table 30: External disk metrics

Metric	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 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.
Read RT (ms)	The average time it took the external disk to serve one read command.
Write RT (ms)	The average time it took the external disk to serve one write command.
Response Time (ms)	The average response time for the reads and writes.
MBs/sec	The size of the IO from the host to the external disk per second.
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.
Avg Queue Depth	Calculated value: Accumulated queue depth/total SCSI command per sec.
% Idle	The percent of time the disk is idle.
% Busy	The percent of time the disk is busy.

External Disk Group metrics

The following table lists all metrics that are available for external disk groups.

Table 31: 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 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.
Read RT (ms)	The average time it took the disk group to serve one read command.
Write RT (ms)	The average time it took the disk group to serve one write command.
Response Time (ms)	
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.
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.
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.

FAST VP Policy metrics

The following table lists all metrics that are available for FAST VP policies.

Table 32: FAST VP policy metrics

Metric	Description
Allocated SG Capacity (GB)	The number of GBs of the storage group that are allocated to the FAST VP 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 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 VP policy.
% Used Capacity (GB)	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.

FE Director metrics

The following table lists all metrics that are available for FE directors.

Table 33: FE director metrics

Metric	Description
% Busy	The percent of time the directory is busy.
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.
Read RT (ms)	The calculated average response time for reads.
Write RT (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.)
Miss Reqs/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.

Metric	Description
% 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.
% 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.
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 0-9	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 0-9	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 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.
% Busy Logical Core 0	
Syscall Time/sec	
Avg RDF Write Time (ms)	
Avg WP Disconnect Time (ms)	

Metric	Description
Avg Read Miss Time (ms)	
Random MBs Written to non WP Slots/sec	
Random Writes to non WP Slots/sec	
Random MBs Written to WP Slots/sec	
Random Writes to WP Slots/sec	
% Busy Logical Core 1	
% Write Req Hit	The percent of write requests that were satisfied from cache.
Read RT Count Range 0 to Read RT Count Range Over 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 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.
Read RT 0 to 1 to Read RT Over 64	These buckets show the distribution of the total average read response time.
Write RT 0 to 1 to Write RT Over 64	These buckets show the distribution of the total average write response time.
Host IO Limit IOs/sec	
Host IO Limit MBs/sec	
Queue Depth Utilization	

FE Emulation metrics

The following table lists all metrics that are available for FE emulations.

Table 34: FE emulation metrics

Metric	Description
Total Read Time	
Total Write Time	
Syscalls Time	
Management Time	
Link Port Time	
XRC Replication Time	
ZHPF Reads Time	
ZHPF Writes Time	
ODX VAAI Xcopy Time	
ODX VAAI Write Same Time	
ODX VAAI Lock Time	
CDI Time	
LP Tasks Time	
SYMMK Total Time	
ORS Time	
SYMMK Idle Time	
Total Work Time	
% Busy	

FE Port metrics

The following table lists all metrics that are available for FE ports.

Table 35: FE port metrics

Metric	Description
Reads/sec	
Writes/sec	
Host IOs/sec	The number of IOs the port is performing on behalf of the specific SG.
MBs Read/sec	
MBs Written/sec	
Host MBs/sec	The number of MBs the port is processing on behalf of the specific SG.
Avg IO Size (KB)	
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	
Read RT (ms)	
Write RT (ms)	
Response Time (ms)	

FE Port - FE metrics

The following table lists all metrics that are available for FE ports (FE).

Table 36: FE port (FE) metrics

Metric	Description
Host IO Limit IOs/sec	
Host IO Limit MBs/sec	
Host IO Limit Exceeded Secs	
Host IO Limit % Time Exceeded	

FE Port - SG metrics

The following table lists all metrics that are available for FE ports (SG).

Table 37: FE port (SG) metrics

Metric	Description
Host IO Limit IOs/sec	
Host IO Limit MBs/sec	
Host IO Limit Exceeded Secs	
Host IO Limit % Time Exceeded	

FICON Emulation metrics

The following table lists all metrics that are available for FICON emulations.

Table 38: FICON emulation metrics

Metric	Description
SYMMK Idle Time	
Total Work Time	
% Busy	
Random Read Hits Time	
Random Read Misses Time	
Random Write Hits Time	
Random Write Misses Time	
Total Sequential Reads Time	
Total Sequential Writes Time	
Syscall Time/secs	
Management Time	
Link Port Time	
XRC Replication Time	
ZHPF Reads Time	
ZHPF Writes Time	
CDI Time	
LP Task Time	
SYMMK Total Time	

Host metrics

The following table lists all metrics that are available for hosts.

Table 39: Host metrics

Metric	Description
Host IOs/sec	
Host Reads/sec	
Host Writes/sec	
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Read RT (ms)	
Write RT (ms)	
Syscall Count/sec	
Host MBs Read/sec	
Host MBs Written/sec	
Response Time (ms)	

Host IO Limit by FE metrics

The following table lists all metrics that are available for host IO limits (by FE).

Table 40: Host IO limit (by FE) metrics

Metric	Description
Host IO Limit IOs/sec	The number of IOs the FE port is processing 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 Exceeded Secs	The number of IOs the FE port is performing on behalf of the specific SG.
Host IO Limit % Time Exceeded	

Host IO Limit by SG metrics

The following table lists all metrics that are available for host IO limits (by SG).

Table 41: Host IO limit (by SG) metrics

Metric	Description
Host IO Limit IOs/sec	The number of IOs the FE port is processing 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 Exceeded Secs	The number of IOs the FE port is performing on behalf of the specific SG.
Host IO Limit % Time Exceeded	

IM Director metrics

The following table lists all metrics that are available for IM directors.

Table 42: IM director metrics

Metric	Description
% Busy	
Memory Page Usage	
Memory Page Fall thru Time	

IM Emulation metrics

The following table lists all metrics that are available for IM emulations.

Table 43: IM emulations metrics

Metric	Description
Page In Time	
Page Out Time	
Environment Time	
Other Time	
SYMMK Total Time	
SYMMK Idle Time	
Total Work Time	
% Busy	

Initiator metrics

The following table lists all metrics that are available for initiators.

Table 44: Initiator metrics

Metrics	Description
Host IOs/sec	
Host Reads/sec	
Host Writes/sec	
Host MBs/sec	
Read RT (ms)	
Write RT (ms)	
Syscall Count/sec	
Host MBs Read/sec	
Host MBs Written/sec	
Response Time (ms)	

Initiators by Port metrics

The following table lists all metrics that are available for initiators (by port).

Table 45: Initiators (by port) metrics

Metrics	Description
Host IOs/sec	
Host Reads/sec	
Host Writes/sec	
Host MBs/sec	
Read RT (ms)	
Write RT (ms)	
Syscall Count/sec	
Host MBs Read/sec	
Host MBs Written/sec	
Response Time (ms)	

IP Interface metrics

The following table lists all metrics that are available for IP interfaces.

Table 46: IP interface metrics

Metric	Description
Packets/sec	
MBs/sec	

iSCSI Target metrics

The following table lists all metrics that are available for iSCSI targets.

Table 47: iSCSI target metrics

Metric	Description
Packets/sec	
MBs/sec	

Masking View metrics

The following table lists all metrics that are available for masking views.

Table 48: Masking view metrics

Metric	Description
Host IOs/sec	
Host MBs/sec	
Response Time (ms)	
Read RT (ms)	
Write RT (ms)	
Capacity (GB)	

Metas metrics

The following table lists all metrics that are available for metas.

Table 49: Metas metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the 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

Metric	Description
	performed by the disk directors to the cache.
Parity Generation Reads	
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
Parity Generation Kbytes Read	
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to the cache.
Read RT (ms)	The average time it took the system to serve one read IO for this volume.
Write RT (ms)	The average time it took the system to serve one write IO for this volume.
Read Miss RT (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 RT (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 RT (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

Metric	Description
	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.
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.

Metric	Description
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 * (total seq IOs per sec / total IOs per sec)
% Seq Read	Calculated value: 100 * (seq reads per sec / 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*(seq writes per sec / 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 * (random read hits per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec/total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
% Random Write Miss	Calculated value: 100 * (random write misses per sec / 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

Metric	Description
	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 RT (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.
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
Random Hits/sec	The total number of hits per second that were not sequential.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
% 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.
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.
Number of RDF Invalid Tracks	
RDF Response Time (ms)	

Metric	Description
IO Density	
BE Partial Sector Writes (KB)	Internal metric.
BE Optimize Writes (KB)	Internal metric.
BE XOR Reads	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE Reads For Copy	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Reads For VLUN Migration	Internal metric.
Reads for Rebuild	
Kbytes Read for Rebuild	
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For VLUN Migration	Internal metric.
Writes for IVTOC	
BE Writes For VLUN Migration (KB)	
Kbytes Written for IVTOC	
BE Writes For Rebuild	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	
Total Capacity (GB)	
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

Metric	Description
	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 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 RT (ms)	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Read Misses/sec	

Metric	Description
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	

Other - Pool Bound Volume metrics

The following table lists all metrics that are available for pool-bound volumes.

Table 50: Pool-bound volumes metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the 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

Metric	Description
	performed by the disk directors to the cache.
Parity Generation Reads	
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
Parity Generation Kbytes Read	
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to the cache.
Read RT (ms)	The average time it took the system to serve one read IO for this volume.
Write RT (ms)	The average time it took the system to serve one write IO for this volume.
Read Miss RT (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 RT (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 RT (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

Metric	Description
	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.
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.

Metric	Description
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 * (total seq IOs per sec / total IOs per sec)
% Seq Read	Calculated value: 100 * (seq reads per sec / 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*(seq writes per sec / 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 * (random read hits per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec/total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
% Random Write Miss	Calculated value: 100 * (random write misses per sec / 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

Metric	Description
	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 RT (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.
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
Random Hits/sec	The total number of hits per second that were not sequential.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
% 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.
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.
Number of RDF Invalid Tracks	
RDF Response Time (ms)	

Metric	Description
IO Density	
BE Partial Sector Writes (KB)	Internal metric.
BE Optimize Writes (KB)	Internal metric.
BE XOR Reads	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE Reads For Copy	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Reads For VLUN Migration	Internal metric.
Reads for Rebuild	
Kbytes Read for Rebuild	
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For VLUN Migration	Internal metric.
Writes for IVTOC	
BE Writes For VLUN Migration (KB)	
Kbytes Written for IVTOC	
BE Writes For Rebuild	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	
Total Capacity (GB)	
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

Metric	Description
	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 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 RT (ms)	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Read Misses/sec	

Metric	Description
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	

Pool by Storage Group metrics

The following table lists all metrics that are available for pools by storage group.

Table 51: Pool by storage group metrics

Metric	Description
BE Reqs/sec	
Allocated Capacity	
IO Density	
BE Read Reqs/sec	
Destage Writes/sec	
BE MBs Transferred/sec	
BE MBs Read/sec	
Destage Write MBs/sec	
Prefetched Tracks/sec	
Pretetched MBs/sec	
Allocated Tracks	
Compressed Tracks	
% Compressed Tracks	
Local Replication Tracks	

Port Group metrics

The following table lists all metrics that are available for port groups.

Table 52: Port group metrics

Metric	Description
Reads/sec	
Writes/sec	
Host IOs/sec	
MBs Read/sec	
MBs Written/sec	
Host MBs/sec	
Avg IO Size (KB)	
% Busy	

RDF Director metrics

The following table lists all metrics that are available for RDF directors.

Table 53: RDF director metrics

Metric	Description
% Busy	The percent of time the directory is busy.
IOs/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 Received/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.
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)
Rewrites/sec	
Write Miss Reqs/sec	
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.
Copy IOs/sec	The number of copy IOs per second via this RDF director. Copy is a generic SRDF term meaning a full track that comes from background data movement. These copies can be generated by a number of different things, including if the RDF device is in adaptive copy mode. Adaptive copy is included in this term, but not exclusively. For example, if SRDF is suspended, hosts continue to write to the R1 and then SRDF is resumed, the resulting R1->R2 copy of tracks

Metric	Description
	(due to invalids) is a “copy.”
Copy MBs/sec	The throughput 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 MSs Sent/sec	The number of async MBs sent per second via this RDF director.

RDF Emulation metrics

The following table lists all metrics that are available for RDF emulations.

Table 54: RDF emulation metrics

Metric	Description
Total Read Time	
Total Write Time	
Total Copy Read Time	
Total Copy Write Time	
Syscalls Time	
Other Time	
LP Task Time	
Link Management Time	
Software Compression Time	
SYMMK Total Time	
SYMMK Idle Time	
Total Work Time	
% Busy	The percent of time that a director is busy.

RDF Port metrics

The following table lists all metrics that are available for RDF ports.

Table 55: RDF port metrics

Metric	Description
Reads/sec	
Writes/sec	
Host IOs/sec	
MBs Read/sec	
MBs Written/sec	
Host MBs/sec	
Avg IO Size (KB)	
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	

RDF/A Group metrics

The following table lists all metrics that are available for RDF/A groups.

Table 56: 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.)
RDF Reads/sec	
Avg Cycle Size	The average size of all the cycles in RDFA sessions.
RDF Writes/sec	
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.

Metric	Description
System WP Limit	The maximum number of write-pending slots for the entire storage 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.
RDF/A WP Count	The number of RDF/A writes pending.
% RDF/A Max Cache Usaged	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.
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).

Metric	Description
RDF/A Session Indicator	Indicates if group is RDFA (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 Number	The active cycle number identifier for the given SRDF/A session. In a regular 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.
Avg IO Service Time (ms)	The average time it took to satisfy reads and writes.
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.
RDF Accumulated Service Time	
Num Devs in Session	The number of storage 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 is 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

Metric	Description
	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 RDF Invalid Tracks	
RDF Response Time (ms)	
Spillover Packed Slots	
Number of Links	The number of links used by this RDF/A group.
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: ((Compressed MBs Received/sec / MBs Received/sec) * 100)
% Compressed MBs Sent	Calculated value: ((Compressed MBs Sent/sec / MBs Sent/sec) * 100)
% Compressed MBs Sent and Received	Calculated value: ((Compressed MBs Sent and Received/sec / MBs Sent and Received/sec) * 100)

Metric	Description
XtremSW Cache Read Hits/sec	
XtremSW Cache Write Hits/sec	
XtremSW Cache Reads/sec	
XtremSW Cache Writes/sec	
XtremSW Cache IOs/sec	
XtremSW Cache Skipped IOs/sec	
XtremSW Cache Dedup Hits/sec	
XtremSW Cache Dedup Reads/sec	
XtremSW Cache Dedup Writes/sec	
% XtremSW Cache Reads	
% XtremSW Cache Writes	
% XtremSW Cache Read Hits	
XtremSW Cache MBs Read/sec	
XtremSW Cache MBs Written/sec	
XtremSW Cache MBs/sec	
XtremSW Cache Read RT (ms)	
XtremSW Cache Write RT (ms)	
XtremSW Cache RT (ms)	
XtremSW Cache Avg Read Size (KB)	
XtremSW Cache Avg Write Size (KB)	
XtremSW Cache Avg IO Size (KB)	
Optimized Read Misses/sec	
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	
RDF Rewrites/sec	
RDF MBs Written/sec	
RDF MBs Read/sec	
RDF Read Hits/sec	

RDF/S Group metrics

The following table lists all metrics that are available for RDF/S groups.

Table 57: RDF/S group metrics

Metric	Description
Host IOs/sec	
Host Reads/sec	
RDF Reads/sec	
RDF Writes/sec	
Host Writes/sec	
Host Hits/sec	
Host Read Hits/sec	
Host Write Hits/sec	
Host Misses/sec	
Host Read Misses/sec	
Host Write Misses/sec	
Host MBs/sec	
Host MBs Read/sec	
Host MBs Written/sec	
BE Reqs/sec	
BE Read Reqs/sec	
BE Write Reqs/sec	
Read RT (ms)	
Write RT (ms)	
Read Miss RT (ms)	
Write Miss RT (ms)	
RDF/S Write RT (ms)	
% Reads	
% Writes	
% Read Hit	
% Write Hit	

Metric	Description
% Read Miss	
% Write Miss	
WP Count	
Seq IOs/sec	
Seq Reads/sec	
Seq Writes/sec	
Seq Read Hits/sec	
Seq Read Miss/sec	
Seq Write Hits/sec	
Seq Write Misses/sec	
Random IOs/sec	
Random Reads/sec	
Random Writes/sec	
Random Read Hits/sec	
Random Write Hits/sec	
Random Read Misses/sec	
Random Write Misses/sec	
Avg IO Size (KB)	
Avg Read Size (KB)	
Avg Write Size (KB)	
% Sequential IO	
% Seq Read	
% Seq Read Hit	
% Seq Read Miss	
% Seq Writes	
% Seq Write Hit	
% Seq Write Miss	
% Random IO	
% Random Write Miss	

Metric	Description
% Random Read Hit	
% Random Read Miss	
% Random Write Hit	
Max WP Threshold	
BE MBs Transferred/sec	
BE MBs Read/sec	
BE MBs Written/sec	
BE Prefetched Tracks/sec	
BE Prefetched Tracks Used/sec	
BE Read Request Time (ms)	
BE Disk Read RT (ms)	
BE Read Task Time (ms)	
% Hit	
% Miss	
% Random Reads	
% Random Writes	
BE % Reads	
BE % Writes	
RDF Response Time (ms)	
Response Time (ms)	
IO Density	
BE Partial Sector Writes (KB)	
BE Optimize Writes (KB)	
BE XOR Reads	
BE XOR Read (KB)	
BE Reads For Copy	
BE Reads For Copy (KB)	
BE Writes For Copy	
BE Writes For Copy (KB)	

Metric	Description
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 (MB)	
Total Capacity (GB)	
Device Block Size	
BE Prefetched MBs/sec	
XtremSW Cache Read Hits/sec	
XtremSW Cache Write Hits/sec	
XtremSW Cache Reads/sec	
XtremSW Cache Writes/sec	
XtremSW Cache IOs/sec	
XtremSW Cache Skipped IOs/sec	
XtremSW Cache Dedup Hits/sec	
XtremSW Cache Dedup Reads/sec	
XtremSW Cache Dedup Writes/sec	
% XtremSW Cache Reads	
% XtremSW Cache Writes	
% XtremSW Cache Read Hits	
XtremSW Cache MBs Read/sec	
XtremSW Cache MBs Written/sec	
XtremSW Cache MBs/sec	
XtremSW Cache Read RT (ms)	
XtremSW Cache Write RT (ms)	
XtremSW Cache RT (ms)	

Metric	Description
XtremSW Cache Avg Read Size (KB)	
XtremSW Cache Avg Write Size (KB)	
XtremSW Cache Avg IO Size (KB)	
Optimized Read Misses/sec	
Optimized MBs Read Miss/sec	
RDF MBs Read/sec	
RDF Read Hits/sec	
RDF MBs Written/sec	
Minimum Echo Delay	
RDF Rewrites/sec	
Avg Optimized Read Miss Size (KB)	
Spillover Packed Slots	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	

SAVE Volume metrics

The following table lists all metrics that are available for SAVE volumes.

Table 58: SAVE volume metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the 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

Metric	Description
	performed by the disk directors to the cache.
Parity Generation Reads	
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
Parity Generation Kbytes Read	
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to the cache.
Read RT (ms)	The average time it took the system to serve one read IO for this volume.
Write RT (ms)	The average time it took the system to serve one write IO for this volume.
Read Miss RT (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 RT (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 RT (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

Metric	Description
	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.
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.

Metric	Description
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 * (total seq IOs per sec / total IOs per sec)
% Seq Read	Calculated value: 100 * (seq reads per sec / 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*(seq writes per sec / 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 * (random read hits per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec/total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
% Random Write Miss	Calculated value: 100 * (random write misses per sec / 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

Metric	Description
	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 RT (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.
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
Random Hits/sec	The total number of hits per second that were not sequential.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
% 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.
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.
Number of RDF Invalid Tracks	
RDF Response Time (ms)	

Metric	Description
IO Density	
BE Partial Sector Writes (KB)	Internal metric.
BE Optimize Writes (KB)	Internal metric.
BE XOR Reads	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE Reads For Copy	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Reads For VLUN Migration	Internal metric.
Reads for Rebuild	
Kbytes Read for Rebuild	
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For VLUN Migration	Internal metric.
Writes for IVTOC	
BE Writes For VLUN Migration (KB)	
Kbytes Written for IVTOC	
BE Writes For Rebuild	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	
Total Capacity (GB)	
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

Metric	Description
	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 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 RT (ms)	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Read Misses/sec	

Metric	Description
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	

Snap Pool metrics

The following table lists all metrics that are available for snap pools.

Table 59: Snap pool metrics

Metric	Description
BE Reqs/sec	
BE Read Reqs/sec	
BE Write Reqs/sec	
BE MBs Transferred/sec	
BE MBs Read/sec	
BE MBs Written/sec	
BE Read Request Time (ms)	
BE Disk Read RT (ms)	
% Used Capacity	
Total Pool Capacity (GB)	
Enabled Pool Capacity (GB)	
Used Pool Capacity (GB)	
Allocated Pool Capacity (GB)	

Spare Disk metrics

The following table lists all metrics that are available for spare disks.

Table 60: Spare disk 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.
IOs/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.
Response Time (ms)	The average response time for the reads and writes.
Read RT (ms)	The average time it took the disk to serve one read command.
Write RT (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.

Metric	Description
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 (GB)	The total capacity of the disk (GBs).
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.

SRP metrics

The following table lists all metrics that are available for SRPs.

Table 61: SRP metrics

Metric	Description
Host IOs/sec	
Host Reads/sec	
Host Writes/sec	
Host MBs/sec	
Host MBs Read/sec	
Host MBs Written/sec	
BE Reqs/sec	
BE Read Req/sec	
BE Write Req/sec	
BE MBs Transferred/sec	
BE MBs Read/sec	
BE MBs Written/sec	
Response Time (ms)	
Read RT (ms)	
Write RT (ms)	

Storage Group metrics

The following table lists all metrics that are available for storage groups.

Table 62: Storage group metrics

Metric	Description
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 IO Limit Exceeded Secs	
Host Hits/sec	The number of host read/write operations performed each second by the group that were immediately satisfied from cache.
Host IO Limit MBs/sec	
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 IO Limit IOs/sec	
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.

Metric	Description
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
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.
Read RT (ms)	The average time that it took the storage to serve one read IO for this group.
Host IO Limit % Time Exceeded	
Write RT (ms)	The average time that it took the storage to serve one write IO for this group.
Read Miss RT (ms)	The average time that it took the storage to serve one read miss IO for this group.
Write Miss RT (ms)	The average time that it took the storage to serve one write miss IO for this group.
RDF/S Write RT (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.

Metric	Description
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.
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

Metric	Description
	sequential.
% Seq Read	Calculated value: 100 * (seq reads per sec / 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*(seq writes per sec / 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 * (random read hits per sec / total IOs per sec)
% Random Write Miss	Calculated value: 100 * (random write misses per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec / total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
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

Metric	Description
	by the disk directors to the cache.
BE Disk Read RT (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.
Read RT Count 5 to 8	
Write RT Count 1 to 3	
Read RT Count Over 32	
Write RT Count 0 to 1	
Read RT Count 14 to 32	
Skew	
Write RT Count 8 to 14	
Write RT Count 14 to 32	
Write RT Count Over 32	
Write RT Count 3 to 5	
Read RT 14 to 32	
Write RT Count 5 to 8	
Read RT Count 8 to 14	
Read RT Over 32	
Write RT 14 to 32	
RDF Response Time (ms)	

Metric	Description
Write RT Over 32	
Read RT Count 0 to 1	
Response Time (ms)	
Read RT Count 1 to 3	
Read RT Count 3 to 5	
IO Density	The number of BE requests per GB of disk.
BE Partial Sector Writes (KB)	The KBs of back-end IO that were partial sector writes.
BE Optimize Writes (KB)	Internal metric.
BE XOR Reads	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE Reads For Copy	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Reads For VLUN Migration	Internal metric.
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE Writes For VLUN Migration (KB)	Internal metric.
BE Writes For Rebuild	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	
Total Capacity (GB)	
Device Block Size	Internal metric.
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.

Metric	Description
Compressed Tracks	The number of compressed tracks for the storage group.
% 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 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 RT (ms)	The average time it took the XtremSW Cache to serve one read.

Metric	Description
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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 Read Misses/sec	Number of read requests each second performed directly from disks bypassing the cache.
Optimized MBs Read Miss/sec	Number of host megabytes read each second directly from disks bypassing the cache.
Avg Optimized Read Miss Size (KB)	
Write paced delay	
Avg. Write Paced Delay (ms)	
Sync Write Reqs/sec	
Sync MBs Sent/sec	
Async Write Reqs/sec	
Async MBs Sent/sec	
Information Alert Count	
Warning Alert Count	
Critical Alert Count	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	

Storage Group by Pool metrics

The following table lists all metrics that are available for storage groups (by pool).

Table 63: Storage group (by pool) 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 Read Reqs/sec	The number of read requests each second performed by the disk directors to cache.
Destage Writes/sec	
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.
Destage Write MBs/sec	
Prefetched Tracks/sec	The number of tracks prefetched from disk to cache in a second.
Pretetched 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.
Local Replication Tracks	

Storage Group by Tier metrics

The following table lists all metrics that are available for storage groups (by tier).

Table 64: Storage group (by tier) 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 Reqs/sec	The number of read requests each second performed by the disk directors to cache.
Destage Writes/sec	
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.
Destage Write MBs/sec	
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.
Allocated Capacity	The total allocated SG capacity.
Compressed Tracks	The number of compressed tracks for the storage group.
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks.
Local Replication Tracks	

Thin Pool metrics

The following table lists all metrics that are available for thin pools.

Table 65: 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.
WP Count	
Avg IO Size (KB)	
Avg Read Size (KB)	
Avg Write Size (KB)	
% Used Capacity	The percent of the pools capacity that is used.
Total Pool Capacity (GB)	The total pool capacity in GBs.
Enabled Pool Capacity (GB)	The enabled pool capacity in GBs.
Used Pool Capacity (GB)	The used pool capacity in GBs.
Allocated Pool Capacity (GB)	The allocated pool capacity in GBs.
BE Read RT (ms)	A calculated value of the response time for all back-end read requests.
BE Write RT (ms)	A calculated value of the response time for all back-end write requests.
BE 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.

Metric	Description
Ingress Tracks	The number of tracks entering the pool.
Egress Tracks	The number of tracks leaving the pool.
BE Pretetched MBs/sec	
Redirect On Write	
Copy on Write	
BE Partial Sector Writes (KB)	
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 Optimize Writes (KB)	
BE XOR Reads	
BE XOR Reads (KB)	
BE Reads For Copy	
BE Reads For Copy (KB)	
BE Writes For Copy (KB)	
BE Blocks Written For Copy (KB)	
BE Blocks Read For VLUN Migration	
BE Reads For VLUN Migration	
BE Writes For VLUN Migration	
BE Blocks Written For VLUN Migration	
BE Writes For Rebuild	
BE Prefetched Tracks/sec	
BE Writes For Rebuild (KB)	
BE RDF Copy	
BE Prefetched Tracks Used/sec	
BE RDF Copy (MB)	
Allocated Capacity (GB)	

Metric	Description
Total Capacity (GB)	
Device Block Size	
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.
Random Write Hits/sec	

Thin Volume metrics

The following table lists all metrics that are available for thin volumes.

Table 66: Thin volume metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the volume.
Host Reads/sec	The total number of host read IO operations performed each second by the volume.
Host Writes/sec	The total number of host write IO operations performed each second by the volume.
Host Hits/sec	The total number of host read IO and write IO operations performed each second by the volume that were immediately satisfied by cache.
Host Read Hits/sec	The total number of host read IO operations performed each second by the volume that were immediately satisfied by cache.
Host Write Hits/sec	The total number of host write IO operations performed each second by the 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 volume.
Host Read Misses/sec	The total number of host read IO operations missed each second by the volume.
Host Write Misses/sec	The total number of host write IO operations missed each second by the 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	
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 RT (ms)	The average time it took the system to serve one read IO for this volume.

Metric	Description
Write RT (ms)	The average time it took the system to serve one write IO for this volume.
Read Miss RT (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 RT (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 RT (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 volume, that were immediately satisfied by cache.
% Write Hit	The percentage of write operations, performed by the volume, that were immediately satisfied by cache.
% Read Miss	The percentage 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.
% Write Miss	The percent of write miss operations performed each second by the 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 volume that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed

Metric	Description
	each second by the volume that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the volume that were misses.
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 * (total seq ios per sec / total IOs per sec)
% Seq Read	Calculated value: 100 * (seq reads per sec / 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*(seq writes per sec / total IOs per sec)
% Seq Write Hit	The percent of the sequential write operations that

Metric	Description
	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 * (random read hits per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec/total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
% Random Write Miss	Calculated value: 100 * (random write misses per sec / total IOs per sec)
Max WP Threshold	The maximum number of write-pending slots available for the volume.
BE MBs Transferred/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 Pretetched 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 Disk Read RT (ms)	
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
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.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
% 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

Metric	Description
	cache to disks.
BE Prefetched MBs/sec	
Number of RDF Invalid Tracks	
RDF Response Time (ms)	
BE Writes for Rebuild	
BE RDF Copy	
BE RDF Copy (MB)	
Allocated Capacity (GB)	
Total Capacity (GB)	
Device Block Size	
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 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.

Metric	Description
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 RT (ms)	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Read Misses/sec	
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	

Tier by Storage Group metrics

The following table lists all metrics that are available for tiers (by storage group).

Table 67: 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 Reqs/sec	The number of read requests each second performed by the disk directors to cache.
Destage Writes/sec	
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.
Destage Write MBs/sec	
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.
Allocated Capacity	The total allocated SG capacity.
Compressed Tracks	The number of compressed tracks for the storage group.
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks.
Local Replication Tracks	

Virtual Pool Tier metrics

The following table lists all metrics that are available for virtual pool tiers.

Table 68: 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 Reqs/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Reqs/sec	
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	
Read RT (ms)	The average time it took the system to serve one read IO for this volume.
Write RT (ms)	The average time it took the system to serve one write IO for this volume.
% Reads	The percentage of IO operations that were reads.
% Writes	The percentage of IO operations that were writes.

Metric	Description

Volume metrics

The following table lists all metrics that are available for volumes.

Table 69: Volume metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the 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.

Metric	Description
Parity Generation Reads	
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
Parity Generation Kbytes Read	
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to the cache.
Read RT (ms)	The average time it took the system to serve one read IO for this volume.
Write RT (ms)	The average time it took the system to serve one write IO for this volume.
Read Miss RT (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 RT (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 RT (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

Metric	Description
	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.
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)

Metric	Description
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 * (total seq IOs per sec / total IOs per sec)
% Seq Read	Calculated value: 100 * (seq reads per sec / 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*(seq writes per sec / 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 * (random read hits per sec / total IOs per sec)
% Random Read Miss	Calculated value: 100 * (random read misses per sec/total IOs per sec)
% Random Write Hit	Calculated value: 100 * (random write hits per sec / total IOs per sec)
% Random Write Miss	Calculated value: 100 * (random write misses per sec / 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

Metric	Description
	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 RT (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.
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
Random Hits/sec	The total number of hits per second that were not sequential.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
% 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.
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.
Number of RDF Invalid Tracks	
RDF Response Time (ms)	
IO Density	
BE Partial Sector Writes (KB)	Internal metric.
BE Optimize Writes (KB)	Internal metric.

Metric	Description
BE XOR Reads	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE Reads For Copy	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Reads For VLUN Migration	Internal metric.
Reads for Rebuild	
Kbytes Read for Rebuild	
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For VLUN Migration	Internal metric.
Writes for IVTOC	
BE Writes For VLUN Migration (KB)	
Kbytes Written for IVTOC	
BE Writes For Rebuild	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	
Total Capacity (GB)	
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.

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 RT (ms)	The average time it took the XtremSW Cache to serve one read.
XtremSW Cache Write RT (ms)	The average time it took the XtremSW Cache to serve one write.
XtremSW Cache RT (ms)	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.
XtremSW Cache Avg IO Size (KB)	The average size of an IO served by XtremSW Cache.
Optimized Read Misses/sec	
Optimized MBs Read Miss/sec	
Avg Optimized Read Miss Size (KB)	

Metric	Description
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	

CHAPTER 8
Databases

This chapter contains the following:

Understanding Database Storage Analyzer 786
Viewing databases 787
Viewing database details 788
Launching Database Storage Analyzer 789

Understanding Database Storage Analyzer

Database Storage Analyzer provides a database troubleshooting solution for Oracle and SQL server databases on storage systems running Enginuity 5671 - 5876 or HYPERMAX OS 5977 or higher.

Database Storage Analyzer supports database to storage correlation by providing a shared view on how performance issues correlate to database level activity and storage level activity. This view is viewable by a database analyst (DBA) and a storage administrator (SA). The view presents I/O metrics such as IOPS, throughput and response time from both the database and the storage system which immediately highlights any gap between the database I/O performance and the storage I/O performance.

In addition, DSA supports hinting capabilities for Oracle databases on storage systems running HYPERMAX OS 5977 or higher that allows users to prioritize mission-critical database processes in order to achieve improved response time. To perform hinting operations, you must be a DSA Admin with hinting permission.

The Database Storage Analyzer help system contains more information.

Viewing databases

This procedure explains how to view database instance information provided by Database Storage Analyzer. StorageAdmin and Database Administrators can use this view to troubleshoot database and storage issues.

Before you begin:

- This feature requires Enginuity 5671 - 5876 or HYPERMAX OS 5977 or higher.
- Only databases monitored by DSA will appear in this view. For instructions on adding monitored databases, refer to the DSA online help.
- Databases associated with redo logs will have two entries in this view, one for the data file, and one for the redo log (indicated with **(REDO)** after the ID).

To view databases:

1. Select the storage system.
2. Select **Databases › Databases** to open the **Database** list view.

   The following properties display:

   - **DB ID** — Database ID. **(REDO)** indicates that the entry is for a redo log.
   - **Environment Name** — Database name.
   - **DB Host** — Name of the host on which the database is installed.
   - **Type** — Type of database application.
   - **Volumes** — Total number of volumes in the storage group used by the database.
   - **Database Capacity (GB)** — Size of the database.
   - **Storage Capacity (GB)** — Total capacity of the database volumes.

   The following control is available:

   - **Database Storage Analyzer** —Launching Database Storage Analyzer on page 789
   - **View Details** — Viewing database details on the next page

Viewing database details

This procedure explains how to view database details.

Before you begin:

This feature requires Enginuity 5671 - 5876 or HYPERMAX OS 5977 or higher.

To view databases:

1. Select the storage system.
2. Select **Databases › Databases** to open the **Database** list view.
3. Select the database and click **View Details** to open its **Details** view.

   The following properties display:

   - **Environment Name** — Database name.
   - **DB Host** — Name of the host on which the database is installed.
   - **Type** — Type of database application.
   - **Number of Volumes** — Total number of volumes in the storage group used by the database.
   - **Storage Capacity** — Total capacity of the database volumes.
   - **Database Capacity** — Size of the database.

   The following control is available:

   - **Database Storage Analyzer** —Launching Database Storage Analyzer on the facing page

Launching Database Storage Analyzer

This procedure explains how to launch the Database Storage Analyzer application.

Before you begin:

This feature requires Enginuity 5671 - 5876 or HYPERMAX OS 5977 or higher.

To launch Database Storage Analyzer:

1. Select the storage system.
2. Select **Databases › Databases** to open the **Database** list view.
3. Click **Database Storage Analyzer**.

CHAPTER 9
System Management

This chapter contains the following:

Monitoring storage systems 792
Viewing storage system details 795
Setting system attributes 799
Setting CPU I/O resource distribution 802
Understanding eNAS 803
Audit log 819
Hardware management 822
Managing jobs 848
Understanding licenses 857
Understanding access controls 862
Understanding dynamic cache partitioning 874
iSCSI 879

Monitoring storage systems

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

Before you begin:

- To access this dashboard, you must be a Monitor, StorageAdmin or higher.
- The information displayed in this dashboard is static, in that it reflects the system's status at the moment you opened the dashboard. To refresh the data, click refresh in the status bar.

To access the System Dashboard:

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

System Dashboard:

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

Summary

- Array nice name, or if not defined, the array serial number. To set a custom name for the array, select the current name, type the new name, and click **OK**. Array names must be unique from other array name and cannot exceed 32 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Array names are case-insensitive.
- Alert summary icon displaying the alert status and count. Alert status is indicated by icon color:

  One or more fatal alerts

  One or more critical alerts, with none higher

  One or more warning alerts, with none higher

  One or more informational alerts, with none higher

  No alerts

  To view details on the system's alerts, click the icon to open the system's alert list view.
- **ID** — Storage system serial number.
- **Model** — Storage system model number.
- **HYPERMAX OS/Enginuity** — Storage system operating environment. When below the minimum suggested version, this value will be in red.
- **Connection** — Whether the storage system is local or remote.
- **Cache Size** — Amount of cache configured on the storage system.

- **D@RE Enabled** — Only displays when Data at Rest Encryption is enabled on the storage system.
- **File Dashboard** — Opens the **File System** dashboard. This option is only available for storage systems running HYPERMAX OS 5977 or higher with file systems installed.
- **Symmetrix Details** — Opens the **Symmetrix Details view**, from which you can view additional details on the storage system.
- **View Reservations** — Opens the **Reservations** list view, from which you can view and manage the system's reservations.

Hardware

Displays the alert status and count for the system's **Front End**, **RDF**, and **Back End** directors, **Available Ports** (HYPERMAX OS 5977 or higher), and **Cache Partitions** (Enginuity 5773 - 5876). Alert status is indicated by icon color:

No alerts and status is normal.

One or more warning alerts, with none higher.

One or more critical alerts, with none higher.

 There are no directors of this particular type on the system

To view additional information on a particular object, click its icon to open the corresponding list view.

If it appears that a director is too busy, click **Available Ports** to open the **Available Ports** list view, from which you can associate additional ports to the director. For more information, refer to Associating directors and ports on page 825.

Health and Maintenance

- **Health Check** — Displays the status of the last health check run on the storage system:

  Not run

  Running

  OK

  Requires attention

  For more information, click the icon to open the **Health and Maintenance** page, from which you can view details on the last health check.

- **Spare Drive Replacement** — (displayed to Administrators or Storage Admins) .Displays the state of the spare drive replacement feature for the storage system:

  Enabled

Disabled

- **Symmetrix Hardware** — Displays the overall state of the Symmetrix hardware:

  OK

  Requires attention

Capacity

Graphic representation of the storage system's capacity. The information displayed depends on the storage operating environment:

HYPERMAX OS 5977 or higher:

- **Virtual** — Displays the sum of free (blue) over total for all thin pools on the storage system.
- **File Storage** — If NAS file systems are available, this chart displays the NAS file system objects free over the total capacity.

Enginuity 5773-5876:

- **Physical** — Displays the sum of free (blue) over total for all the disks on the storage system.
- **Virtual** — Displays the sum of free (blue) over total for all thin pools on the storage system.

Other Hardware

Displays alert status and count for other types of directors, if they are present on the system. Possible types include, External, IM, and EDS. Alert status is indicated by icon color:

No alerts and status is normal.

One or more warning alerts, with none higher.

One or more critical alerts, with none higher.

Viewing storage system details

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **Symmetrix Details** to open the **Symmetrix Details** view.

   The following properties display:

   - **Array Display Name** —Custom *nice name* assigned to storage system. if any. To assign a name, or change an existing name, click **Add Name** to open the **Set Array Display Name** dialog box. Array names must be unique from other array names and cannot exceed 32 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Array names are case-insensitive.
   - **Product Model**
   - **Microcode Version (Number)**
   - **Microcode Date/Patch Date**
   - **Capacity**
   - **Used Capacity**—not applicable to storage systems running HYPERMAX OS 5977.
   - **Free capacity**— not applicable to storage systems running HYPERMAX OS 5977.
   - **Actual Capacity**—not applicable to storage systems running HYPERMAX OS 5977.
   - **Cache Size (Mirrored)**
   - **Number of Front End Directors**
   - **Number of Back End Directors**
   - **Number of RDF Directors**
   - **Number of Configured (Symm) Devices**
   - **Number of Visible (Host) Devices**
   - **Number of Configured Actual Disks**
   - **Number of Configured Spare Disks**
   - **Number of Unconfigured Disks**
   - **Max Number of Hypers per Disk**
   - **Device Masking (VCM) Config State**
   - **VCMDB Access Restricted State**
   - **Device Masking (ACLX) Config State**
   - **ACLX Volume LUN Address**
   - **Access Control Configuration State**
   - **Configuration Change State**

- **Disk Group Assignments**
- **Parity RAID Configuration**
- **RAID Configuration**
- **PAV Mode**
- **PAV Alias Limit**
- **SDDF Configuration State**
- **WORM Configuration Level**
- **WORM Characteristics**
- **Switched RDF Configuration State**
- **Concurrent RDF Configuration State**
- **Dynamic RDF configuration State**
- **Concurrent Dynamic RDF State**
- **RDF Data Mobility Configuration State**
- **SRDF/A Maximum Host Throttle (Secs)**
- **SRDF/A Maximum Cache Usage (Percent)**
- **Multi LRU Device Assignment**
- **# of Available Cache Slots**
- **Max # of DA Write Pending Slots**
- **Max # of System Write Pending Slots**
- **Max # of Device Write Pending Slots**
- **Max # of Replication Slots**—not applicable to storage systems running HYPERMAX OS 5977.
- **Symmetrix Last IPL Time (Cold)**
- **Symmetrix Last Fast IPT Time (Hot)**
- **Symmetrix Alerts (Enabled/Disabled)**
- **Max Capacity of DSE reservation**—not applicable to storage systems runningHYPERMAX OS 5977.
- **Symmetrix Priority Controls**—not applicable to storage systems running HYPERMAX OS 5977.
- **Hot Swap Policy**
- **Symmetrix Disk Library**

- **FBA Geometry Emulation**
- **3 Dynamic Mirrors**
- **Cache Partitioning**
- **Auto Meta**
- **Minimum Auto Meta Size (Cyl)**
- **Auto Meta Member Size (Cyl)**—not applicable to storage systems running HYPERMAX OS 5977.
- **Auto Meta Configuration**
- **Disk Service State**
- **Front door LED status**
- **Symmetrix Data Encryption**
- **Replication Cache Usage (Percent)**—not applicable to storage systems running HYPERMAX OS 5977.

Setting system attributes

Before you begin:

Depending the storage system operation environment, some of the following attributes may not apply.

To set system attributes:

1. Select the storage system.
2. Select **System › Settings › Symmetrix Attributes** to open the **Symmetrix Attributes** page.
3. Modify any number of the attributes.

   Attributes are organized into the following panels:
   General:

   - **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 storage systems running Enginuity versions lower than 5874.
   - **VCMDB Access Restricted** —Specify whether the VCM database can be accessed on this storage system. This is only applicable to DMX.

     Select to enable, or clear to disable.

   Enable Auto Meta:

   This panel only displays for storage systems running Enginuity 5773 - 5876.

   - **Enable Auto Meta** — 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 Enginuity 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 Enginuity 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.

PAV:

This panel only displays for storage 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 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.

SRDF/A:

- **Maximum Cache Usage (Percent)** — Type the maximum percentage of system write-pending cache slots for all SRDF/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.

SRDF DIRECTORS:

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.

DCP:

- **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 876.
- **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.

1. 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 799.

1. Select the storage system.
2. Select **System › System Dashboard**.
3. In the **Hardware** panel click **RDF** to open the **RDF** 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.
   - **Director Name** — Type the director name.
   - **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**.

Understanding eNAS

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

With this integrated storage solution, Unisphere for VMAX StorageAdmin provision VMAX storage to eNAS data movers, which triggers the creation of storage pools in VNX. Unisphere for VNX users then use the storage pools for file-level provisioning (for example, creating file systems, file shares, etc.)

Unisphere for VMAX provides the following features to support eNAS:

- **File System dashboard** — Provides a central location from which to monitor and manage integrated VNX file services. For more information, refer to Managing File Storage below.
- **Provision Storage for File wizard** — Allows you to provision VMAX storage to eNAS data movers. For more information, refer to Provisioning storage for file on page 805.
- **Launch Unisphere for VNX** — Allows you to link and launch Unisphere for VNX. For more information, refer to Launching Unisphere for VNX on page 808.

Discovering eNAS control stations

This procedure explains how to connect to an eNAS control station for the first time.

Before you begin:

- To perform this operation, you must be an Administrator.
- The storage system must be running HYPERMAX OS 5977 or higher.

To connect to an eNAS control station:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **Credentials** dialog box.
4. Type the username and password you use to access the eNAS control station and click **OK**.

Managing File Storage

The **File Dashboard** provides you with a single place from which to monitor and manage integrated VNX file services.

Before you begin:

- To perform this operation, you must be a Monitor or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.
- The **File Storage** dashboard provides a point-in-time view of the eNAS instance. To refresh the information displayed in the dashboard, click refresh storage system in the status bar, or exit/return to the dashboard.

To access the File Dashboard:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard**.
4. If prompted, type the username and password you use to access the eNAS control station, and click **OK**.

File Dashboard:

The **File Storage** dashboard is organized into the following panels:

Summary:

- **File Systems** — File systems associated with the file storage groups.
- **File Storage Groups** — File storage groups associated with the file systems.
- **File Masking Views** — Masking views associated with the file storage groups.

Capacity:

Displays the following capacity information:

- **Virtual** — Free versus total capacities of the file storage groups.
- **File** — Free versus total capacities for the file systems associated with the file storage groups on the storage system.

Most Consumed Capacity:

Displays the file storage pools with the most consumed capacity. To view a list of all file storage pools, regardless of their consumed capacity, click **View All** to open the **Storage Pools** list view.

Common Tasks:

Links to common file storage tasks, including:

- Provisioning storage for file on the facing page
- Launching Unisphere for VNX on page 808

Control Stations:

Displays the name and status of the control stations.

Data Movers:

Displays the status of the control station's data movers:

- **Data Mover Name** — Data mover name.
- **Status** — Health status of the data mover.

File Storage Alerts:

Details all alerts associated with the eNAS instance (file systems, storage groups, masking views), including:

- **Severity** — Alert's severity. Possible values are:
  - (1) Fatal
  - (2) Critical
  - (3) Warning
  - (4) Information
  - (5) Normal
- **Description** — Description of the alert.
- **Created** — Data/time the alert was issued.

Provisioning storage for file

This procedure explains how to provision VMAX storage to eNAS datamovers.

Before you begin:

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running HYPERMAX OS 5977 or higher.

To provision storage for file:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Dashboard**.
4. If prompted, type the username and password you use to access the eNAS control station, and click **OK**.
5. In the **Common Tasks** panel, click **Provision Storage for File** to open the **Provision Storage for File** wizard.
6. Type a **Storage Group Name** name. 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.
7. To select a **Storage Resource Pool** other than the default (DEFAULT_SRP), click Edit ( ) and select one. To create the storage group outside of FAST control, select **None**.
8. Select the **Service Level** to set on the SG. Service levels specify the characteristics of the provisioned storage, including average response time, workload type, and priority. This field defaults to None if you set the **Storage Resource Pool** to None. Possible values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critial, high rate OLTP
Gold	High	Very heavy I/O, database logs, datasets
Silver	Price/Performance	Database datasets, virtual applications
Bronze	Cost optimized	Backup, archive, file
Optimized (Default)		Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

9. Refine the service level by selecting the **Workload Type** to assign to it
10. Type the number of **Volumes** and select the **Capacity** of each.
11. *Optional*: Click edit ( ) and set any number of the following advanced options for the storage group:
    - To have FAST factor the R1 volume statistics into move decisions made for the R2 volume, select **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 storage system is part of an SRDF setup.

- To allocate capacity for each SG volume, select **Allocate capacity for each volume**.

12. *Optional*: To add a child storage group, click **Add Service Level** and specify a storage group **Name**, **Service Level**, **Workload Type**, **Volumes** and **Capacity**. Repeat this step for each additional child storage group. The maximum number of allowed child storage groups is 64.
13. *Optional*: To set host I/O limits for the storage groups, click **Set Host I/O Limits** to open the **Host I/O Limits** dialog box. For instructions setting the limits, refer to dialog's help page.
14. Click **Next**.
15. Select a **VNX Host** and click **Next**. Only hosts with at least one initiator visible to the eNAS instance are available.
16. Select whether to use a **New** or an **Existing** port group, and then do the following depending on your selection:

    New:

    a. *Optional*: Edit the suggested **Port Group Name** by highlighting it and typing a new name over it. Port group names must be unique from other port groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Port group names are case-insensitive.

    b. Select the ports to use. To view host-invisible ports (unmasked and unmapped), select **Include ports not visible to the host**.

    The following properties display:

    - **Dir:Port** — Storage system director and port in the port group.
    - **Identifier** — Port identifier.
    - **Initiators Logged In** — Number of initiators logged into the fabric.
    - **PGs** — Number of port groups where the port is a member.
    - **% Busy** — Percentage of time that the port is busy.

    c. Click **Next**.

    Existing: Select the port group and click **Next**. Only port groups with at least one port visible to the eNAS instance are available.

17. *Optional*: Edit the suggested name for the **Masking View** by highlighting it and typing a new name over it. Verify the rest of your selections. To change any of them, click **Back**. Note that some changes may require you to make additional changes to your configuration.
18. *Optional*: Determine if the storage system can handle the updated service level:

    a. Click **Run Suitability Check**. The **Suitability Check** dialog box opens, indicating the suitability of the change. For information on interpreting the results, refer to the dialog's help page. This option is only available under certain circumstances. For more information, refer to Suitability Check restrictions.

    b. Click **OK** to close the message.

    c. If your updates are found to be unsuitable, modify the settings and run the check again until the suitability check passes.

19. *Optional*: To manually set the LUN addresses for the masking operation:

    a. Click **Set Dynamic LUNs** to open the **Set Dynamic LUN Addresses** dialog box.

b. Notice the address displayed in the **Starting LUN** field. To accept this automatically generated address, click **Apply Starting LUN**. To move to the next available, click **Next Available LUN**.

c. When done, click **OK** to close the dialog.

20. Back on the wizard **Review** page, do one of the following:

- 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 851 and Previewing on page 850.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

Launching Unisphere for VNX

This procedure explains how to launch Unisphere for VNX.

Before you begin:

- To perform this operation, you must be a StorageAdmin or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To launch Unisphere for VNX:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. If prompted, type the username and password you use to access the eNAS control station, and click **OK**.
5. In the **Common Tasks** panel, click **Launch Unisphere for VNX**.

Managing file storage groups

This procedure explains how to manage file storage groups.

Before you begin:

- To perform this operation, you must be a StorageAdmin or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To manage file storage groups:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Dashboard**.
4. In the **Summary** panel, click **File Storage Groups** to open the **File Storage Groups** list view.
5. Select one or more storage groups and click one of the following, depending on the operation your want to perform:
   - **Modify** —Modifying storage groups on page 129
   - **Provision Storage for File** —Provisioning storage for file on page 805
   - **View Details** —Viewing storage group details (HYPERMAX OS 5977 or higher) on page 150
   - **FAST Array Advisor** — Analyzing FAST migration on page 202
   - **Change SRP** —Changing Storage Resource Pools for storage groups on page 170
   - **Delete** —Deleting storage groups on page 145
   - **Start Allocate/Free/Reclaim** —Managing thin pool allocations on page 263/Managing thin pool capacity on page 267/Managing space reclamation on page 265
   - **Stop Allocated/Free Reclaim** —Managing thin pool allocations on page 263/Managing thin pool capacity on page 267/Managing space reclamation on page 265
   - **Convert to Cascaded** —Converting storage groups to cascaded on page 140
   - **Set Volume Status**—Setting volume status on page 476

Managing file masking views

This procedure explains how to manage file masking views.

Before you begin:

- To perform this operation, you must be a StorageAdmin or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To manage file masking view:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Dashboard**.
4. In the **Summary** panel, click **File Masking Views** to open the **File Masking Views** list view.
5. Select one or more masking views and click one of the following, depending on the operation your want to perform:
   - **View Connections** — See Viewing masking view connections on page 339.
   - **Delete** — See Deleting masking views on page 339.
   - **View Details** — See Viewing masking view details on page 341.

Viewing file systems

This procedure explains how to view VNX file systems.

Before you begin:

- To perform this operation, you must be a Monitor or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To view file systems:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Dashboard**.
4. In the **Summary** panel, click **File Systems** to open the **File Systems** list view.

   The following properties display:

   - **File System Name** — Name of the file system.
   - **Storage Groups** — Storage groups to which the file system belongs.
   - **Storage Pool** — Storage pool to which the file system belongs. If the file system does not belong to a storage pool, the default name, **Meta Volume**, will display in this column.
   - **Allocated %** — Allocated capacity of the file system in GB.
   - **Capacity (GB)** — Total capacity of the file system in GB.

   The following control is available:

   View Details - Viewing file system details below

Viewing file system details

This procedure explains how to view details on VNX file systems.

Before you begin:

- To perform this operation, you must be a Monitor or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To view file system details:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. In the **Summary** panel, click **File Systems** to open the **File Systems** list view.
5. Select a file system and click **View Details** to open its **Details** view.

   The file system **Details** view contains Properties and Related Objects panels.

Properties

The following properties display:

- **File System Name** — Name of the file system.
- **File System Description** — Description of the file system.
- **File System Unique ID** — Unique VNX ID for the file system.
- **File System Total Capacity** — Size of the file system.
- **File System Free Capacity** — Available space on the file system.
- **File System Block Size** — Size of a block on the file system.
- **File System Type** — Type of file system and its conventions (for example, NTFS).
- **Operational Status** — Current operational status of the local file system.

Related Objects:

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 **Storage Groups - 2** will open a view listing the storage groups associated with the file system.

Viewing file system storage pools

This procedure explains how to view storage pools associated with VNX file systems.

Before you begin:

- To perform this operation, you must be a StorageAdmin or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To view file system storage pools:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. In the **Summary** panel, click **File Systems** to open the **File Systems** list view.
5. Select a file system and click **View Details** to open its **Details** view.
6. In **Related Objects** panel, click **Storage Pools** to open the **Storage Pools** list view.

   The following properties display:

   - **Storage Pool Name** — Name of the storage pool.
   - **Instance ID** — Long name description of the storage pool.
   - **File System** — File system using the storage pool in a VNX context.

   The following controls are available:

   - **View Details** — Viewing file system storage pool details below

Viewing file system storage pool details

This procedure explains how to view storage pools associated with VNX file systems.

Before you begin:

- To perform this operation, you must be a StorageAdmin or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To view file system storage pools:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. In the **Summary** panel, click **File Systems** to open the **File Systems** list view.
5. Select a file system and click **View Details** to open its **Details** view.
6. In **Related Objects** panel, click **Storage Pools** to open the **Storage Pools** list view.
7. Select a storage pool and click **View Details** to open it details view.

   The storage pools details view includes Properties, Related Objects, and File Systems panels.

Properties panel:

The following properties display:

- **Name** — Name of the storage pool.
- **Instance ID** — Long name description of the storage pool.
- **Health State** — Current health of the storage pool. Possible values include:
  - Unknown
  - OK
  - Degraded/Warning
  - Minor Failure
  - Major Failure
  - Non recoverable error
- **Total Managed Space (GB)** — Amount of capacity usable for allocation of storage volumes, logical disks, or child storage pools.
- **Remaining Managed Space (GB)** — Remaining usable capacity after allocation of storage volumes, logical disks, or child storage pools
- **Operational Status** — Status for the operational condition of the pool. Possible values include:
  - Unknown
  - Not Available
  - Servicing
  - Starting
  - Stopping
  - Abandoned

- Dormant
- Completeted
- Migrating
- Emigrating
- Immigrating
- Snaopshotting
- Shutting Down
- In Test
- Transitioning
- In Service

◆ **Pool ID** — Unique ID of the storage pool.

◆ **Primordial** — Whether the storage pool can be created by consumers of this model.

◆ **Usage** — Intended usage or any restrictions on the pool.

Related Objects panel:

The **Related Objects** panel links you to views displaying objects contained in and associated with the storage pool. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Storage Groups - 3** will open a view listing the three storage groups associated with the storage pool.

File Systems panel:

◆ **File System Name** — Name of the file system.

◆ **Storage Groups** — Storage groups mapped to the file system.

◆ **Storage Pools** — Storage pools in the file system.

◆ **Total Capacity** — Total capacity in GB of the file system.

◆ **Free Capacity** — Free capacity in GB of the file system.

◆ **Nr of File** — Number of files.

◆ **Alerts** — File system alerts.

Managing file storage alerts

Viewing file storage alerts

This procedure explains how to view alerts related to file storage.

Before you begin:

- To perform this operation, you must be a Monitor or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To view file storage alerts:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. In the **File Storage Alerts** panel, click **View All** to open the **File Storage Alerts** list 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. In this view, this value will always be File.
   - **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.

   The following controls are available:

   - **View Details** — Viewing file storage alert details below
   - **Acknowledge** — Viewing file system storage pool details on page 813
   - **Delete** — Deleting file storage alerts on the facing page

Viewing file storage alert details

This procedure explains how to view details on alerts related to file storage.

Before you begin:

- To perform this operation, you must be a Monitor or higher.
- The storage system must be running HYPERMAX OS 5977 or higher.

To view file system storage pools:

1. Select the storage system.
2. Select **System › System Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. In the **File Storage Alerts** panel, click **View All** to open the **File Storage Alerts** list view.
5. Select an alert and click View Details to open the file storage **Alert Details** view.

   The following properties display:

   - **Alert ID** — Unique number assigned by 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 storage system generating the alert.
   - **Object** — Object to which the alert is related. For more information, click the object to open its details view.
   - **Created** — Date/time the alert was created.
   - **Description** — Description of the alert.

   The following controls are available:

   - **Acknowledge** —Viewing file system storage pool details on page 813
   - **Delete** —Deleting file storage alerts below

Deleting file storage alerts

You can delete file storage alerts from the overall **Alerts** list view or the **File Storage Alerts** view.

To delete from the overall Alerts list view:

1. Select the storage system.
2. Select **System › Alerts** to open the **Alert** list view.

3. Select one or more **File** type alerts and click **Delete**.

To delete from the File Storage Alerts view:

1. Select the storage system.
2. Select **System › Dashboard** to open the **System** dashboard.
3. In the **Summary** panel, click **File Dashboard** to open the **File Storage** dashboard.
4. In the **File Storage Alerts** panel, click **View All** to open the **File Storage Alerts** list view.
5. Select one or more alerts and click **Delete**.

Audit log

Viewing the system audit log

The storage system audit records come from the SYMAPI database and include all actions taken on that storage system. The audit log resides on the storage 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 system 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 the next page.
- The user must have Auditor or higher permissions.

To view the system audit log:

1. Select the storage 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 storage system.
   - **Username** — Name of the user operating on the storage system.
   - **Host** — Host operating on the storage system.
   - **Function Class** — Generic audit category for the operation on the storage system.
   - **Action Code** — Specific audit code for the operation on the storage 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 storage 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 storage system.
   - **Username** — Name of the user operating on the storage system.
   - **Function Class** — Generic audit category for the operation on the storage system.
   - **Action Code** — Specific audit code for the operation on the storage system.
   - **Host** — Host operating on the storage 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 storage system.
   - **Application Version** — Version of the application operating on the storage 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 specifed 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).

Before you begin:

This procedure requires Enginuity 5671 - 5876.

To convert directors:

1. Select a storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. Depending on the type of director you are converting, do the following in the **Hardware** panel:
   Front end directors:
   a. Click **Front End** to open the **Front End Directors** list view.
   b. Select a director, click more >> and click **Convert FA to RF** to open the **Convert Front End Director to RDF Director** dialog box.
   c. To view additional information about the selected director, expand **Show Selected Directors**.
   d. 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 851 and Previewing on page 850.
      - Expand **Add to Job List**, and click **Run Now** to convert the director now.

   RDF directors:
   a. Click **RDF** 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. To view additional information about the selected director, expand **Show Selected Directors**.
   d. 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 851 and Previewing on page 850.
      - 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.

Associating directors with ports

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click **Available Ports** to open the **Available Ports** list view.
4. Select a port, click **Associate** to open the **Port Association** dialog box.
5. Select a director and click **OK**.

Setting director port attributes

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click the type of director to open its 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 attributes are for all port types. Therefore, depending on the port type, some of the attributes may not be available to you.

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	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.
Non Participating	Non participating.
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.

7. 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 851 and Previewing on page 850.

Associating directors and ports

When associating directors and ports, you can associate directors with ports, or ports with directors.

Before you begin:

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

To associate directors with ports:

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click **Available Ports** to open the **Available Ports** list view.
4. Select one or more ports, click **Associate**to open the **Associate Director** dialog box.
5. Select a director and click **Associate**.

To associate ports with directors:

1. Select the storage system.

2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the appropriate panel (**Hardware** or **Other Hardware**), click the director emulation type to open its list view.
4. Select the director, click more >> and click **Associate** to open the **Port Association** dialog box.
5. Select the port and **OK**.

Disassociating directors and ports

Before you begin:

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

To disassociate ports with directors:

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the appropriate panel (**Hardware** or **Other Hardware**), click the director emulation type to open its list view.
4. Select the director, click more >> and click **Disassociate**.
5. Click **OK**.

Enabling and disabling director ports

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click the type of director to open its list view.
4. Select one or more ports, and click **Enable** to enable them, or **Disable** to disable them.
5. 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 › System Dashboard** to open the system **Dashboard**.
3. In the **Health and Maintenance** panel, 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.

Naming storage systems

This procedure explains how to assign custom *nice names* to storage systems:

1. Select a storage system
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Summary** panel, click **Symmetrix Details** to open the **Symmetrix Details** view.
4. In the **Array Display Name** field, click either the existing name or **Add Name** to open the **Set Array Display Name** dialog box.
5. Type an **Array Display Name**. Array names must be unique from other array names and cannot exceed 32 characters. Only alphanumeric characters, underscores (_), and (-) are allowed. Array names are case-insensitive.
6. Click **OK**.

Locating VMAX systems

The following procedures explain how to locate a 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**. This feature is only available on storage systems running Enginuity 5874 - 5876.
- 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 storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Health and Maintenance** panel, click **Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.
4. In the **Hardware Components** 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 storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Health and Maintenance** panel, click **Symmetrix Hardware** to open the **Symmetrix Hardware** dashboard.
4. 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**.

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 on the facing page

Step 2: Replace the failed drive on the facing page

Step 3: Verify the drive replacement on page 831

Before you begin:

- To perform this procedure you must be an Administrator or Storage Admin.

- The storage system with the failed disk drive must be locally attached to the SMAS server, running Enginuity 5876.163.105 or higher, and have the CRU flag enabled. This feature is not supported on storage systems running HYPERMAX OS 5977 or higher.
- Ensure that there are no other disk replacements occurring on the storage system.
- Ensure that you have the replacement drive and storage system's 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 storage system
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Health and Maintenance** panel, 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 851Scheduling jobs on page 851.

   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 853):

```
Spare Drive Replacement SUCCEEDED
```

If any other messages appear, follow the instructions.

Viewing hardware components

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

From this view you can:

- View system components in a tree view or details view.
- View the current health or state of each system component.
- Control bay door LED blinking to identify a system bay. Locating VMAX systems on page 827 includes instructions on using this feature.

To access the System Hardware dashboard:

1. Select the storage system.
2. Select **System › Dashboard** to open the system **Dashboard**.
3. In the **Health and Maintenance** panel, click **Symmetrix Hardware** to open the **System Hardware** dashboard.

   The dashboard opens with the top-level system configuration as the default view.
4. Drill-down into the **Hardware Components** menu to view status and details on the system's components.

   The health of each component is indicated with a status icon:

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

   In the tree menu, component status is propagated up through the menu to the top system level. To locate a degraded or failed component, drill down to find the faulty component.

Viewing available ports

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click **Available Ports** to open the **Available Ports** list view.

   The **Available Ports** list view allows you to view and associate available ports with directors.

   The following properties are displayed:

   - **Slot** — Slot number.
   - **Port** — Port ID.
   - **Type** — Port type.
   - **Speed** — Transmission rate (input/output channel).

   The following control is available:

   - **Associate** — Associating directors with ports on page 823

Viewing back-end directors

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

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

   Note Some of the following properties and controls are environment-specific and therefore may appear in a different order or not at all.

   The following properties display:

   - **Director** — Director name.
   - **Dir % Busy** — Percent of time the director is busy. This field only displays when the storage system is registered with the Performance component.
   - **Type** — Director type.
   - **Hypers** — Number of hyper volumes serviced by the director. This field only displays for storage systems running Enginuity 5773 - 5876.
   - **Protocol** — Transmission protocol for the director port.
   - **Cores** — Number of director cores.
   - **Director Status** — Director status.
   - **Port Interface** — Port interface ID. This field only displays for storage systems running Enginuity 5773 - 5876.
   - **Port Serviced Hypers** —Number of hyper volumes serviced by the director.
   - **Port Status** — Indicates whether the port is online or offline.

   The following controls are available:

   - **View Details** — See Viewing back-end director details on the next page.

Viewing back-end director details

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click **Back End** 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 Ojects, Performance Views panels.

Properties panel

The following properties display:

- **Director** — 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

Before you begin:

External directors are only supported on storage systems running HYPERMAX OS 5977 or higher.

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Other Hardware** panel (bottom, right), click **External Directors** to open the **External Directors** list view.

   Use this view to view the external directors.

   Note Some of the following properties and controls are environment-specific and therefore may appear in a different order or not at all.

   The following properties display:

   - **Director** — Director name.
   - **Port** — Port ID.
   - **Director Status** — Director status.
   - **Cores** — Number of available director cores.
   - **Hypers** — Number of hyper volumes serviced by the director.
   - **Alerts** — Director alert status, as indicated by icon color:

     One or more fatal alerts

     One or more critical alerts, with none higher

     One or more warning alerts, with none higher

     One or more informational alerts, with none higher

     No alerts

   The following control is available:

   - **View Details** — Viewing external director details on the next page
   - **Associate** — Associating directors and ports on page 825 (This control only displays when the storage system is running HYPERMAX OS 5977 or higher.)
   - **Disassociate** — Disassociating directors and ports on page 826 (This control only displays when the storage system is running HYPERMAX OS 5977 or higher.)

Viewing external director details

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Other Hardware** panel (bottom, right), 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:

- **Director** — 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 storage system is not registered for data collection.

Viewing system front-end directors

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

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

   Note Some of the following properties and controls are environment-specific and therefore may appear in a different order or not at all.

   The following properties display:

   - **Director**— Director name.
   - **Port** — Port number.
   - **Alerts** — Director alert status, as indicated by icon color:

     One or more fatal alerts

     One or more critical alerts, with none higher

     One or more warning alerts, with none higher

     One or more informational alerts, with none higher

     No alerts
   - **Type**—Director transmission protocol.
   - **Director Status** — Whether the director is online or offline.
   - **Port Status** — Indicates whether the port is online or offline.
   - **Port ID** — Port ID.
   - **Volumes** — Number of volumes mapped to the director port.
   - **ACLX** — Indicates whether access control logic is enabled or disabled.
   - **Cores** — Number of cores. This field only displays when the storage system is running HYPERMAX OS 5977 or higher.
   - **Speed GB/Sec**— The speed of the individual cores on the system (applicable only to storage systems running HYPERMAX OS 5977 or higher.)
   - **VCM State** — Indicates whether VCM volume state is 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 826.
- **Disable** — See Enabling and disabling director ports on page 826.
- **Set Port Attributes** — See Setting director port attributes on page 823.
- **Convert FA to RF** — See Converting directors on page 822. (not applicable to storage systems running HYPERMAX OS 5977 or higher.)
- **Set ORS Ceiling** — See Setting Open Replicator ceiling on page 546.
- **Map** — Mapping volumes on page 226 (Only enabled when ACLX is disabled)
- **Unmap** — Unmapping volumes on page 227 (Only enabled when ACLX is disabled)
- **Associate** — Associating directors and ports on page 825 (Only displays when the storage system is running HYPERMAX OS 5977 or higher.)
- **Disassociate** — Disassociating directors and ports on page 826 (Only displays when the storage system is running HYPER MAX OS 5977 or higher.)

Viewing system front-end director details

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click **Front End** 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 to view and manage a front end director. It contains Properties, Related Objects, and Performance Views panels.

Properties panel

The following properties display:

- **Director** — Director name.
- **Port** — Port number.
- **Type**—Director transmission protocol.
- **Director Status** — Director status.
- **Port Status** — Indicates port status, ON or OFF.
- **Director Connection Status** — Connection status.
- **Volumes** — Number of volumes mapped to the port.
- **Mapped Devices (including meta members)** — Number of mapped volumes, including any meta members.
- **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.
- **VNX Attached** — Indicates whether VNX Attached is enable 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.
- **No Participating** — Indicates whether Non Participate feature is enabled or disabled.
- **SPC2 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 Enginuity 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.
- **Sequent** — Indicates whether the Sequent feature is enable or disabled. This feature is used for Sequent platforms.

- **Server on 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.
- **Open Replicator Maximum Ceiling** — Maximum recopy ceiling value.
- **Open Replicator Ceiling (%)** — Recopy ceiling value.
- **Open Replicator Actual Ceiling** — Recopy actual 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 826.
- **Disable** — See Enabling and disabling director ports on page 826.
- **Set Port Attributes** — See Setting director port attributes on page 823.
- **Set ORS Ceiling** — See Setting Open Replicator ceiling on page 546.

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

   Use this view to view and manage the RDF directors.

   Note Some of the following properties and controls are environment-specific and therefore may appear in a different order or not at all.

   The following properties display:

   - **Director** — Director ID
   - **Dir % Busy** — Percent of time the director is busy. This field only displays when the Symmetrix system is registered with the Performance component.
   - **Type**—Director type.
   - **RDF groups** — Number of RDF groups.
   - **Port** — Port ID.
   - **Director Status** — Director status.
   - **RDF Protocol** — Director transmission protocol.
   - **IPSec Policies** — Number of IPSec Policies.
   - **Port Status** — Indicates whether the port is online or offline.
   - **Cores** — Number of cores. This field only displays when the storage system is running HYPERMAX OS 5977 or higher.
   - **Alerts** — Director alert status, as indicated by icon color:

     One or more fatal alerts

     One or more critical alerts, with none higher

     One or more warning alerts, with none higher

     One or more informational alerts, with none higher

     No alerts

The following controls are available:

- **View Details** — See Viewing RDF director details below.
- **Convert RF to FA** — See Converting directors on page 822.(not applicable to storage systems running HYPERMAX OS 5977 or higher.)
- **CPU I/O Setting** — See Setting CPU I/O resource distribution on page 802.
- **Associate** — Associating directors and ports on page 825 (only applicable to storage systems running HYPERMAX OS 5977 or higher.)
- **Disassociate** — Disassociating directors and ports on page 826(only applicable to storage systems running HYPERMAX OS 5977 or higher.)

Viewing RDF director details

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Hardware** panel, click **RDF** 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 this attribute is true or false.
- **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.

Viewing IM directors

Before you begin:

Infrastructure Manager (IM) directors are only supported on storage systems running HYPERMAX OS 5977 or higher.

To view IM directors:

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Other Hardware**panel, click **IM directors** to open the **IM Directors** list view.

   The following properties display:

   Note Some of the following properties are environment-specific and therefore may appear in a different order or not at all.

   - **Director** — Director name.
   - **Director Type** — Director type.
   - **Director Status** — Director status.
   - **Cores** — Number of available director cores.
   - **Associated Ports** — Number of associated ports.
   - **Alerts** — Director alert status, as indicated by icon color:

     One or more fatal alerts

     One or more critical alerts, with none higher

     One or more warning alerts, with none higher

     One or more informational alerts, with none higher

     No alerts

Viewing EDS directors

Before you begin:

Enginuity Data Services (EDS) directors are only supported on Symmetrix systems running Enginuity version 5977 or higher.

To view EDS directors:

1. Select the storage system.
2. Select **System › System Dashboard** to open the system **Dashboard**.
3. In the **Other Hardware** panel, click **EDS directors** to open the **EDS Directors** list view.

   The following properties display:

   Note Some of the following properties are environment-specific and therefore may appear in a different order or not at all.

   - **Director** — Director name.
   - **Dir % Busy** — Percent of time the director is busy. This field only displays when the Symmetrix system is registered with the Performance component.
   - **Director Type** — Director type.
   - **Director Status** — Director status.
   - **Cores** — Number of available director cores. This field only displays when the Symmetrix system is running Enginuity version 5977 or higher.
   - **Associated Ports** — Number of associated ports.
   - **Alerts** — Director alert status, as indicated by icon color:

     One or more fatal alerts

     One or more critical alerts, with none higher

     One or more warning alerts, with none higher

     One or more informational alerts, with none higher

     No alerts

Managing jobs

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

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 70: 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 storage system 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 Customer Service for assistance in reverting to your previous configuration should there be unforeseen problems with the new configuration.

Understanding task persistence

Active 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 71: 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

1. Select the storage system.
2. Select **System › Job List** to open the **Job List** view.

Scheduling jobs

1. Select the storage 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 storage system.
2. Select **System › Job List** to open the **Job List** view.
3. Select a job and click **Run**.
4. Click **OK**.

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 storage 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 storage 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 storage system.

Reordering tasks within a job

1. Select the storage 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 storage 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 storage 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 Unisphere for VMAX.
5. Click **OK**.

Un-grouping jobs

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

1. Select the storage 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 **Ungroup**for each of the tasks jobs you want to un-group from the job.
5. Click **Apply**.

Stopping jobs

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

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

Deleting jobs

1. Select the storage 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.

   Unisphere for VMAX will remove the jobs from the job list view.

Viewing the job list

1. Select the storage 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 storage system.

   The following properties display:

   - **Name** — User-defined name for the job or an ID assigned by 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).
     - **SUCCEEDED**— Job finished running.
     - **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.
     - **FAILED** — 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 851
   - **View Details** — Viewing job details on page 855
   - **Group** — Grouping jobs on the previous page
   - **Stop** — Stopping jobs on the previous page

- **Delete** — Deleting jobs on the previous page

Viewing job details

1. Select the storage 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.

   The job **Details** view allows you to view and manage a job.

   The following properties are displayed:

   - **Name** — User-defined name for the job or an ID assigned by 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.
   - **Scheduled** — Date and time the job is scheduled to run.
   - **Owner** — Host from which the job was created and the ID of user who created it.
   - **Last Modified** — Date and time the job was moved to the job list.
   - **Completed** — Date and time the job completed This field is blank for incomplete or failed jobs.
   - **Task list** — Lists the configuration tasks in the job.
     - **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.
   - **Result** — Displays the config change steps as the job executes, and the results once the job finishes.

   The following controls are available:

   - **Run** — See Running jobs on page 851.
   - **Schedule** — See Scheduling jobs on page 851.

- **Delete** — See Deleting jobs on page 853.
- **Stop** — See Stopping jobs on page 852.
- **Ungroup** — See Un-grouping jobs on page 852.
- **Apply** — Applies changes made in the job. For example, renaming the job.
- **Cancel** — Cancels changes made in the **Properties** list.

Understanding licenses

Unisphere for VMAXsupports Electronic Licensing (eLicensing). 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.

When installing licenses with eLicensing, you obtain license files from EMC, copy them to a Solutions Enabler or a Unisphere for VMAX host, and push them out to storage 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 EMC Online Support and push it out to the storage system.

When managing your licenses, Solutions Enabler, Unisphere for VMAX, 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

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

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

Unisphere for VMAX licensing

Unisphere for VMAX uses array-based eLicensing.

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

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

Note For more information on eLicensing, refer to the *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 Unisphere for VMAX to manage a storage system, you must first install a Unisphere for VMAX license on the system.
- To obtain array-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 storage system.
2. Select **System › Settings › Electronic Licenses** to open the **Electronic License** page.
3. Do the following, depending on the license type:

   Array-based licenses:

   a. Click **Symmetrix Entitlements** to open the **Symmetrix Entitlements** list view.
   b. Click **Get File** to open the 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 **OK**.

Removing host-based licenses

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

Viewing Symmetrix entitlements

1. Select the storage system.
2. Select **System › Settings › Electronic Licenses › Symmetrix Entitlements** to open the **Symmetrix Entitlements** list view.
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 array-based entitlements and host-based keys, which are still required) on storage systems running Enginuity 5875 or higher. In addition, this view also allows you obtain and install license files on storage systems running Enginuity 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 storage systems or to all the storage 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 array-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. This value only applies to storage systems running Enginuity 5875 - 5876.
  - **R-TB-SATA** — Indicates that the capacity licensed applies to the raw capacity of all SATA volumes on the system. his value only applies to storage systems running Enginuity 5875 - 5876.
  - **REG-TB** — Indicates that the capacity licensed applies to the registered capacity of the storage system. his value only applies to storage systems running Enginuity 5875 - 5876.
  - **R-TB External** — Indicates that the capacity licensed applies to the raw capacity of the virtualized LUNs in external storage. his value only applies to storage systems running Enginuity 5875 - 5876.
  - **Usable-TB** — Indicates that the capacity licensed applies to the usable capacity of the storage system. This value only applies to storage systems running HYPERMAX OS 5977 or higher.
- **Capacity (TB)**— 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 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 Symmetix system.
- **Load File** — See Installing licenses on page 857.

Viewing host-based licenses

1. Select the storage 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 857.
   - **Remove Key** —See Removing host-based licenses on page 858.

Viewing license usage

1. Select the storage system.
2. Select **System › Settings › Electronic Licenses › License Usage** to open the **License Usage** view.

   The **License Usage** view shows the results of a query to the storage 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 Override**— 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 (TB)** — Maximum quantity of data for which the functionality of the software is licensed to use, in Terabytes.
- **Used Capacity (TB)** — Amount of licensed capacity currently in use.

Understanding access controls

Administrators , **StorageAdmins**, and **SecurityAdmins** can set access controls on specific volumes within a storage system and assign those volumes to a specific host. When set, only that host can see the volumes, and perform the granted operations. Other hosts connected to that storage 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 See Creating access groups below..
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 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 See Creating access pools on page 864.
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 See Creating access control entries on page 866..

Refer to the *Solutions Enabler 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 storage 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 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 storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Access 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 Enginuity 5771 and lower) or 31 characters (Enginuity 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 storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** and click **OK**.
4. Click **Access 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 storage system.
2. Select **System › Settings › Symmetrix Access Controls** .
3. Type the **Access Control Pin** and click **OK**..
4. Click **Access 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 storage system.
2. Select **System › Settings › Symmetrix Access Controls** .
3. Type the **Access Control Pin** and click **OK**.
4. Click **Access 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 storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access Groups** to open the **Access Groups** list view.

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

   The following properties display:

   - **Access ID Group** — 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 862.
   - **View Details** — See Viewing access group details on page 869.
   - **Delete** — See Deleting access groups above.

Creating access pools

1. Select the storage 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 names must be unique from other access pools on the system and cannot exceed 8 characters Enginuity 5771 and lower) or 31 characters (Enginuity 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 more 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 storage 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 storage 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 storage 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 storage 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 Control Entries** — Number of access groups associated with the pool.

   The following controls are available:

   - **Create** — See Creating access pools on page 864.
   - **View Details** — See Viewing access pool details on page 870.
   - **Delete** — See Deleting access pools on the previous page.

Modifying access types

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

1. Select the storage 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 See Access types on page 872..

   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 storage system.
2. Select **System › Settings › Symmetrix Access Controls**.

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

Viewing access control entries

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

   Use the **ACEs** list view to display and manage access control entries on the storage 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 872..

   The following controls are available:

   - **Create** — See Creating access control entries on the previous page.
   - **View Details** — See Viewing access control entry details below.
   - **Delete** — See Deleting access control entries above.

Viewing access control entry details

1. Select the storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.

4. Click **ACEs** 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 872..

The following controls are available:

- **Create** — See Creating access control entries on page 866.
- **Delete** — See Deleting access control entries on the 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 storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access Groups** to open the **Access Groups** list view.

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

   The following properties display:

   - **Access ID Group** — 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 862.
   - **View Details** — See Viewing access group details on the facing page.
   - **Delete** — See Deleting access groups on page 864.

Viewing access group details

1. Select the storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **Access 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 ID Group** — 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 862.
- **Delete** — See Deleting access groups on page 864.

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 storage 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 863.

- **Remove** — See Removing access IDs from access groups on page 864.

Viewing access pools

1. Select the storage 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 storage 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 Control Entries** — Number of access groups associated with the pool.

   The following controls are available:

   - **Create** — See Creating access pools on page 864.
   - **View Details** — See Viewing access pool details below.
   - **Delete** — See Deleting access pools on page 865.

Viewing access pool details

1. Select the storage 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 864.
- **Delete** — See Deleting access pools on page 865.

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 storage 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 865.
   - **Remove** — See Modifying access pools on page 865.

Viewing access types

1. Select the storage system.
2. Select **System › Settings › Symmetrix Access Controls**.
3. Type the **Access Control Pin** or select **Read Only** and click **OK**.
4. Click **ACEs** 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 the next page...

   The following controls are available:

   Add — Adding access ID to access groups on page 863

   - **Remove** — Removing access IDs from access groups on page 864

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 storage 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.

Understanding dynamic cache partitioning

Dynamic Cache Partitioning (DCP) divides the cache memory into multiple partitions with unique names and their device path assignments. Partition areas can be made static or dynamic in size. The dynamic partitioning provides flexibility to the amount of floating memory that can be allocated with a high and low watermark. This allows memory resources to be temporarily donated to other partitions when needed. The symqos command allows you to create partitions for different device groupings in addition to the default partition that all devices belong to initially. Each partition has a target cache percentage as well as a minimum and maximum percentage. In addition, you can donate unused cache to other partitions after a a specified donation time.

Note This feature requires Enginuity 5771 - 5876.

Enabling/Disabling dynamic cache partitioning

Before you begin:

This feature is not supported on HYPERMAX OS 5977 or higher.

To enable/disable dynamic cache partitioning:

1. Select a storage 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 storage system is defined in the storage 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 storage system.
2. Select **System › System Dashboard**.
3. In the **Hardware** panel click **Cache Partitions** to open the **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 storage 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 dynamic cache partitions

This procedure explains how to assign dynamic cache partitions from the **Volumes** view. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

1. Select the storage system.
2. Select **Storage › Volumes**.
3. Locate the volumes using the volumes filter (HYPERMAX OS 5977 or higher), or the tree menu in the **Volumes** dashboard (Enginuity 5671 - 5876). For information on using these components, refer to Managing volumes (HYPERMAX OS 5977 or higher) on page 208 and Managing volumes (Enginuity 5671 - 5876) on page 208, respectively.
4. Select the volumes, click more >> , and select **Assign Dynamic Cache Partition** to open the **Assign Dynamic Cache Partition** dialog box.
5. *Optional*: Click **Show Selected Volumes** to view details on the selected volumes.
6. Select a **Dynamic Cache Partition** and click **OK**.
7. Click **OK** in the confirmation message.

Deleting dynamic cache partitions

1. Select the storage system.
2. Select **System › System Dashboard**.
3. In the **Hardware** panel, click **Cache Partitions** to open the **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 874.
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 storage system.
2. Select **System › System Dashboard**.
3. In the **Hardware** panel, click **Cache Partitions** to open the **Cache Partitions** list view.

   Use the **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.
- **Alerts** — Director alert status, as indicated by icon color:

  One or more fatal alerts

  One or more critical alerts, with none higher

  One or more warning alerts, with none higher

  One or more informational alerts, with none higher

  No alerts

The following controls are available:

- **Create** — See Creating dynamic cache partitions on page 874.
- **View Details** — See Viewing dynamic cache partition details below.
- **Delete** — See Deleting dynamic cache partitions on page 875.

Viewing dynamic cache partition details

1. Select the storage system.
2. Select **System › System Dashboard**.
3. In the **Hardware** panel, click **Cache Partitions** to open the **Cache Partitions** list view.
4. Select a partition and click **View Details** to open it **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 storage 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 values 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 (Sec)**— 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 cached used by the partition.

The following controls are available:

- **Create** — See Creating dynamic cache partitions on page 874.
- **Delete** — See Deleting dynamic cache partitions on page 875.
- **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 will display with inactive links if the selected storage system is not registered for data collection.

iSCSI

Viewing the iSCSI dashboard

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

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

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

The iSCSI dashboard provides you with a graphical view of the relationship between an iSCSI director and its associated iSCSI ports, iSCSI targets, IP interfaces and IP routes. From the dashboard, you can navigate to monitor and manage iSCSI directors, iSCSI ports, iSCSI targets, IP interfaces and IP routes.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.

   The following panels display:

   - **iSCSI DIRECTORS**—A panel
     - displaying the overall number of and relationships between the iSCSI objects - iSCSI Directors, unattached iSCSI Targets, attached iSCSI Targets, iSCSI Ports, IP Routes, and IP Interfaces. Clicking on an item opens the associated list view.
     - listing the iSCSI directors. The status of the iSCSI director is represented by an icon that is based on the alerts for the director. Clicking **All Items**results in the relationship panel displaying all e iSCSI Directors. Selecting a specific director results in an update to relationship panel to display director information for the selected director.
   - **iSCSI Management** — A panel with the following controls:
     - **Create iSCSI Target** — Creating an iSCSI target on page 883
     - **Create IP Interface** — Creating an IP interface on page 885
     - **Add IP Route** — Adding an IP route on page 886
   - **iSCSI Alerts**— A panel listing the latest alerts listing the associated object, description and creation date/time. Click **View All** to view all alerts.

     The following properties display:
     - **State** — The state of the alert.
     - **Severity** — The severity of the alert.
     - **Type** — The alert type.

- **Object** — The object associated with the alert.
- **Description** — The description of the alert.
- **Acknowledged**— The date and time that the alert was acknowledged.

Creating an iSCSI target

This procedure explains how to create an an iSCSI target.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached) to open the (Attached or Unattached) **iSCSI Targets** list view.
4. Click **Create**.
5. Select a director from the drop-down menu.
6. Optional: Type a value for the **Target Name**. If you do not enter a value, the system generates a unique iqn for the target name.

   If you enter a value, it must start with 'iqn.' or 'eui.' strings, be composed of alphanumeric characters, colons, dashes, and periods, not be longer than 255 characters, and be unique on the storage system.
7. Type a value for the **Network ID**.Valid values range from 1 through 16383.
8. Type a value for the **TCP Port**. Default value is 3260. Valid values range from 0 through 65535.
9. Optional: Click **Show Advanced** to view the port properties advanced section for the iSCSI target.

   a. Select one or more of the following **Port Flags**:

   - **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 (the storage systems do not currently support the upper range of volume set).
   - **Environ Set** — Enables the environmental error reporting by the storage system to the host on the specific port.
   - **Avoid Reset Broadcast** — Enables a SCSI bus reset to only occur to the port that received the reset (not broadcast to all channels).
   - **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** — 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.
   - **Open VMS** — Enables an Open VMS fiber connection.

- **ISID Protected** — Protected initiator session identifier. This flag applies to iSCSI target only. It has no effect for FA/FCoE ports. When this flag is set on iSCSI target ,the system uses the ISI) and iSCSI name to calculate the World Wide Name (WWN) in the Persistent Group Registration (PGR) record.

10. Click **OK**.

Creating an IP interface

This procedure explains how to create an IP interface.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Interface** to open the **IP Interfaces** list view.
4. Click **Create**.
5. Select a director and port combination from the drop-down menu.
6. Type a value for the **IP Address**.You specify the IPv4 address in dotted decimal notation. You specify the IPv6 address in colon-hexadecimal format. The IP address must be unique within a SE Director Emulation/Network ID combination.
7. Type a value for the **Prefix**. You specify an IPv4 prefix length value in the range of 1-30. You specify an IPv6 prefix length value in the range of 1-128.
8. Type the value for the **Network id**.Valid values range from 1 through 16383.
9. Type a value for the **Vlan id**. Valid values range from 0 through 4094. This setting will fail if same VLAN id is used for more than one IP interface on a specified SE physical port.
10. Optional: Type a value for the **Max Transmission Unit**. Default value is 1500. Valid values range from 1500 through 9000.
11. Click **OK**.

Adding an IP route

This procedure explains how to add an IP route.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › iSCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Route** to open the **IP Routes** list view.
4. Click **Add IP Route**.
5. Select a director and port combination from the drop-down menu.
6. Type a value for the **IP Address**.You specify the IPv4 address in dotted decimal notation. You specify the IPv6 address in colon-hexadecimal format.
7. Type a value for the **Prefix**. You specify an IPv4 prefix length value in the range of 1-30. You specify an IPv6 prefix length value in the range of 1-128.
8. Type a value for the **Gateway IP**.You specify an IPv4 or IPv6 address. You specify the IPv4 address in dotted decimal notation. You specify the IPv6 address in colon-hexadecimal format.
9. Optional: Type the value for the **Network ID**. Default value is 1. Valid values range from 1 through 16383.
10. Click **OK**.

Deleting an iSCSI target

This procedure explains how to delete an iSCSI target. Deleting iSCSI Targets with IP Interfaces attached will detach the IP Interfaces before deleting the iSCSI Target.

The deletion fails if the target is in a port group.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › iSCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached) to open the **iSCSI Targets** list view.
4. Select an iSCSI target and click **Delete**.
5. Click **OK**.

Deleting an IP interface

This procedure explains how to delete an IP interface.

Note: An IP interface attached to an iSCSI target cannot be deleted.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › iSCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Interfaces** to open the **IP Interfaces** list view.
4. Select an IP interface and click **Delete**.
5. Click **OK** to confirm the operation.

Removing an IP Route

This procedure explains how to remove an IP route.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › iSCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Routes** to open the **IP Routes** list view.
4. Select an IP route and click **Remove**.
5. Click **OK** to confirm the operation.

Attaching an IP interface to an iSCSI target

This procedure explains how to attach an IP interface to an iSCSI target.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Interface** to open the **IP Interfaces** list view.
4. Select an IP Interface and click **Attach**.
5. Select whether to attach to an existing or new iSCSI target, and do the following depending on your selection.

   Existing: Select an iSCSI target from the list filtered by network ID and director.

   New:

   a. Click **Use Custom Name** and type a value for the **Target Name**.If you do not enter a value, the system generates a unique iqn for the target name.

      If you enter a value, it must start with 'iqn.' or 'eui.' strings, be composed of alphanumeric characters, colons, dashes, and periods, not be longer than 255 characters, and be unique on the storage system.

   b. Optional: Type a value for **TCP Port**. Default value is 3260. Valid values range from 0 through 65535.

   c. Optional: Click **Show Advanced** to display the port flag properties for the target being created.

      Select one or more of the following **Port Flags**:

   - **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 (the storage systems do not currently support the upper range of volume set).
   - **Environ Set** — Enables the environmental error reporting by the storage system to the host on the specific port.
   - **Avoid Reset Broadcast** — Enables a SCSI bus reset to only occur to the port that received the reset (not broadcast to all channels).
   - **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** — 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.

- **Open VMS** — Enables an Open VMS fiber connection.
- **ISID Protected** — Protected initiator session identifier. This flag applies to iSCSI target only. It has no effect for FA/FCoE ports. When this flag is set on iSCSI target ,the system uses the ISI) and iSCSI name to calculate the World Wide Name (WWN) in the Persistent Group Registration (PGR) record.

6. Click **OK**. Please note that if you toggle between the **New** and **Existing** selections, the system only applies the changes from within your last selection.
7. Click **OK** to confirm your changes.

Attaching an iSCSI target to an IP interface

This procedure explains how to attach an iSCSI target to an IP interface.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached) to open the **iSCSI Targets** list view.
4. Select an iSCSI target and click **Attach**.
5. Select whether to use an existing or new IP interface(s), and do the following depending on your selection.

   Existing: Select an IP interface from the list. Note: You can selected a maximum of 8 interfaces (the total (8) includes the number of already attached interfaces).

   New:

   a. Select a director and port combination from the drop-down menu.
   b. Type a value for the **IP Address**.You specify the IPv4 address in dotted decimal notation. You specify the IPv6 address in colon-hexadecimal format.
   c. Type a value for the **Prefix**.You specify an IPv4 prefix length value in the range of 1-30. You specify an IPv6 prefix length value in the range of 1-128.
   d. Type a value for the **VLAN ID**. Valid values range from 0 through 4094. This setting will fail if same VLAN id is used for more than one IP interface on a specified SE physical port.
   e. View the value for the **Network ID**. Value displayed is based on the existing target selection.
   f. Optional: Type a value for the **Max Transmission Unit**. Default value is 1500. Valid values range from 1500 through 9000.
   g. Optional: Click **Add IP Interface** to add one or more IP interfaces.
6. Click **OK**. Please note that if you toggle between the **New** and **Existing** selections, the system only applies the changes from within your last selection.
7. Click **OK** to confirm your changes.

Detaching an IP interface from an iSCSI target

This procedure explains how to detach an IP interface from an iSCSI target.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Interfaces** to open the **IP Interfaces** list view.
4. Select an IP Interface and click **Detach**.
5. Click **OK** to confirm the operation.

Disabling an iSCSI target

This procedure explains how to disable an iSCSI target.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** to open the (Attached or Unattached) **iSCSI Targets** list view.
4. Select an iSCSI target and click **Disable**.
5. Click **OK**.

Enabling an iSCSI target

This procedure explains how to enable an iSCSI target.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** to open the (Attached or Unattached) **iSCSI Targets** list view.
4. Select an iSCSI target and click **Enable**.
5. Click **OK**.

Setting port flags

This procedure explains how to set port flags.

The iSCSI target must be in an offline state before the flags can be modified.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached) to open the **iSCSI Targets** list view.
4. Select an iSCSI target and click **Set Port Flags**.
5. Select one or more of the following **Port Flags**:

- **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 (the storage systems do not currently support the upper range of volume set).
- **Environ Set** — Enables the environmental error reporting by the storage system to the host on the specific port.
- **Avoid Reset Broadcast** — Enables a SCSI bus reset to only occur to the port that received the reset (not broadcast to all channels).
- **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** — 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.
- **Open VMS** — Enables an Open VMS fiber connection.
- **ISID Protected** — Protected initiator session identifier. This flag applies to iSCSI target only. It has no effect for FA/FCoE ports. When this flag is set on iSCSI target ,the system uses the ISI) and iSCSI name to calculate the World Wide Name (WWN) in the Persistent Group Registration (PGR) record.

6. Click **OK**.

Viewing the iSCSI directors list

This procedure explains how to view the iSCSI directors list.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › iSCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Directors** to open the **iSCSI Directors** list view.

   The following properties display, depending on the operating environment:

   - **Director**— The name of the iSCSI director.
   - **Ports**— The number of iSCSI physical ports associated with the director.
   - **IP Interfaces** — The total number of IP interfaces associated with each physical port on the director.
   - **Attached iSCSI targets**—The total number of unique iSCSI targets attached to the IP interfaces associated with the director.
   - **Unattached iSCSI targets**— The total number of unique iSCSI targets unattached to the IP interfaces associated with the director .
   - **IP Routes**— The total number of IP routes associated with the director.
   - **Volumes** — The total number of volumes mapped to all iSCSI targets on that director.

   The following controls are available, depending on the operating environment:

   - **Create Target**— See Creating an iSCSI target on page 883.
   - **Add IP Route** — Adding an IP route on page 886
   - **View Details**— See "Viewing the iSCSI director details" on page 894

Viewing the iSCSI director details

This procedure explains how to view the iSCSI director details.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Directors** to open the **iSCSI Directors** list view.
4. Select a row from the list and click **View Details** to open the **iSCSI Directors**details view.

Properties panel

The following properties display, depending on the operating environment:

- **Director**— The name of the iSCSI director.
- **Number of iSCSI Ports**— The number of the iSCSI physical port associated with the director. (Should Ports be Port ?)
- **Number of IP Interfaces** — The total number of IP interfaces associated with each physical port on the director.
- **Number of Attached iSCSI Targets** —The total number of iSCSI targets attached to the IP interfaces associated with the director.
- **Number of Unattached iSCSI targets** — The total number of iSCSI targets unattached to the IP interface sassociated with the director. .
- **Number of IP Routes**— The total number of IP routes associated with the director.
- **Volumes** — The total number of volumes mapped to all iSCSI targets on that port.

Related Objects panel

The following objects display, depending on the operating environment:

- **iSCSI Ports**— The list view of iSCSI ports that are associated with this iSCSI director.
- **IP Interfaces** — The list view of IP interfaces that are associated with the physical ports on this director.
- **Attached iSCSI Targets** —The list view of iSCSI targets on this director that are attached to IP interfaces.
- **Unattached iSCSI Targets** — The list view of iSCSI targets on this director that are unattached.
- **IP Routes**— The list view of IP routes associated with this iSCSI director.
- **Mapped Volumes** — The list view of volumes that are mapped to iSCSI targets associated with the director.

Viewing the IP interfaces list

This procedure explains how to view the IP interfaces list.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Interfaces** to open the **IP Interfaces** list view and select a specific director to filter the list.

   The following properties display, depending on the operating environment:

   - **Dir:Port**— The combination of the physical director and associated port of the IP interface.
   - **Network ID**—The network identity of the IP interface that provides isolated IP routes.
   - **IP Address**— The destination IPv4 or IPv6 address of the IP interface.
   - **Prefix**— The network mask IPv4 or IPv6 prefix value of the IP interface.
   - **VLAN ID** — The VLAN tag number of the IP interface.
   - **Dir:Virtual Port**— The combination of the physical director and assigned virtual port or iSCSI target alias of the IP interface.
   - **Volumes** — The total number of volumes mapped to the iSCSI target attached to the IP interface.

   The following controls are available, depending on the operating environment:

   - **Create**—See Creating an IP interface on page 885.
   - **Attach**— See "Attaching an iSCSI target to an IP interface" on page 890
   - **Detach**— See "Detaching an IP interface from an iSCSI target" on page 891
   - **Delete**—See "Deleting an IP interface" on page 887
   - **View Details**— See "Viewing the IP interfaces details" on page 896

Viewing the IP interfaces details

This procedure explains how to view IP interfaces details.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Interface** to open the **IP Interface** list view and select a specific director to filter the list.
4. Select a row from the list and click **View Details** to open the **IP Interfaces** details view.

Properties panel

The following properties display, depending on the operating environment:

- **Dir:Port**— The combination of the physical director and the associated port of the IP interface.
- **Network ID**—The network identity of the IP interface that provides isolated IP routes. This can also be modified.
- **IP Address**— The destination IPv4 or IPv6 address of the IP interface. This can also be modified.
- **Prefix**— IPv4 or IPv6 prefix length. This can also be modified. You specify an IPv4 prefix length value in the range of 1-30. You specify an IPv6 prefix length value in the range of 1-128.
- **VLAN ID**— The VLAN tag number of the IP interface.
- **Maximum Transmission Unit** — The maximum transit size of the ethernet packet for this IP interface. This can also be modified.
- **Dir:Virtual Port**— The combination of the physical director and the assigned virtual port or iSCSI target alias of the IP interface.
- **iSCSI target**— The iSCSI target IQN for the attached target.
- **Volumes**— The total number of volumes mapped to the iSCSI target attached to the IP interface.
- **Volume Set Addressing**
- **Environ Set**
- **Avoid Reset Broadcast**
- **Disable Q Reset on UA**
- **Soft Reset**
- **SCSI 3**
- **SCSI Support1(OS2007)**
- **SPC2 Protocol Version**
- **Open VMS**
- **ISID Protected**

The following controls are available, depending on the operating environment:

- **Delete**— See Deleting an IP interface on page 887.
- **Apply**— Apply updates.
- **Cancel**— Cancel updates.

Related Objects panel

The following objects display, depending on the operating environment:

- **iSCSI Targets** —The list view of iSCSI targets listing the single attached iSCSI target.
- **iSCSI Ports**— The list view of iSCSI ports listing the single attached iSCSI port.
- **Mapped Volumes** — The list view of volumes that are mapped to iSCSI targets associated with the director.

Viewing the iSCSI targets list

This procedure explains how to view the iSCSI targets list.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached) to open the **iSCSI Targets** list view and select a specific director to filter the list.

   The following properties display, depending on the operating environment:

   - **Name**— The IQN of the iSCSI target.
   - **Dir:Virtual Port**— The assigned virtual port for the iSCSI target.
   - **Status**— The status of the iSCSI target.
   - **Network ID**—The network identity of the IP interface that provides isolated IP routes.
   - **IP Interfaces**— The total number of IP interfaces attached to the iSCSI target.
   - **iSCSI Ports**— The total number of physical IP ports associated with the iSCSI target.
   - **Volumes** — The total number of volumes mapped to the iSCSI target.

   The following controls are available, depending on the operating environment:

   - **Create** — See Creating an iSCSI target on page 883.
   - **Enable**— See "Enabling an iSCSI target" on page 891
   - **Disable**— See "Disabling an iSCSI target" on page 891
   - **Set Port Flags**— See "Setting port flags" on page 892
   - **Attach**— See "Attaching an iSCSI target to an IP interface" on page 890
   - **Detach**— Detaches the selected iSCSI target from the IP interface.
   - **View Details**— See "Viewing the iSCSI target details" on page 900
   - **Delete**— See "Deleting an iSCSI target" on page 886

Viewing the iSCSI target details

This procedure explains how to view the iSCSI target details.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached) to open the **iSCSI Targets** list view and select a specific director to filter the list.
4. Select a row from the list and click **View Details** to open the **iSCSI Targets** details view.

Properties panel

The following properties display, depending on the operating environment:

- **Name**— The IQN of the iSCSI target (editable).
- **Dir:Virtual Port**— The assigned virtual port for the iSCSI target.
- **Network ID**—The network identity of the IP interface that provides isolated IP routes (editable).
- **TCP**— The TCP port to be used for all IP addresses attached to it (editable).
- **Status** — The status of the iSCSI target.
- **Number of IP Interfaces**— The total number of IP interfaces attached to the iSCSI target.
- **Number of iSCSI Ports**— The total number of physical iSCSI ports associated with the iSCSI target.
- **Volumes** — The total number of volumes mapped to the iSCSI target.

The following controls are available, depending on the operating environment:

- **Enable**— See Enabling an iSCSI target on page 891.
- **Disable**— See Disabling an iSCSI target on page 891.
- **Set Port Flags**— See "Setting port flags" on page 892
- **Delete**— See "Deleting an iSCSI target" on page 886
- **Apply**— Apply updates.
- **Cancel**— Cancel updates.

Related Objects panel

The following objects display, depending on the operating environment:

- **iSCSI Ports**— The list view of iSCSI ports that are associated with this iSCSI target.
- **IP Interfaces** — The list view of IP interfaces that are associated with the physical ports on this iSCSI target.
- **Volumes** — The list view of volumes that are mapped to this iSCSI target.

Viewing the IP routes list

This procedure explains how to view the IP routes list.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Routes** to open the **IP Routes** list view and select a specific director to filter the list.

   The following properties display, depending on the operating environment:

   - **Destination IP**— The IPv4 or IPv6 destination IP address of the IP route.
   - **Prefix**— The IPv4 or IPv6 network mask prefix of the IP route.
   - **Gateway IP**— The gateway IP address of the IP route.
   - **Network ID**—The network identity of the IP route.
   - **Director**— The director that the IP route is associated with.

   The following controls are available, depending on the operating environment:

   - **Add IP Route** — Adding an IP route on page 886
   - **Remove**— See Removing an IP Route on page 887.
   - **View Details**— See "Viewing the IP routes details" on page 902

Viewing the IP routes details

This procedure explains how to view the IP routes details.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **IP Routes** to open the **IP Routes** list view and select a specific director to filter the list.
4. Select a row from the list and click **View Details** to open the **IP Routes** details view.

Properties panel

The following properties display, depending on the operating environment:

- **Destination IP**— The IPv4 or IPv6 destination IP address of the IP route.
- **Prefix**— The IPv4 or IPv6 network mask prefix of the IP route.
- **Gateway IP**— The gateway IPv4 or IPv6 address of the IP route.
- **Network ID**—The network identity of the IP route.
- **Director**— The director that IP route is associated with.

The following control is available, depending on the operating environment:

- **Remove**— See Removing an IP Route on page 887.

Related Objects panel

The following objects display, depending on the operating environment:

- **iSCSI Director** —The list view containing the associated director for the selected IP route.

Viewing the iSCSI ports list

This procedure explains how to view the iSCSI ports list.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Ports** to open the **iSCSI Ports** list view and select a specific director to filter the list.

   The following properties display, depending on the operating environment:

   - **Director**— The name of the iSCSI director.
   - **Port**— The iSCSI physical port number associated with the director.
   - **Status**— The current status of the physical port (offline or online).
   - **IP Interfaces** — The total number of IP interfaces associated with the physical port.
   - **iSCSI targets**—The number of iSCSI targets that are attached to to the IP interfaces.
   - **Volumes** — The total number of volumes mapped to all iSCSI targets on that port.

   The following controls are available, depending on the operating environment:

   - **Create IP Interface**— See Creating an IP interface on page 885.
   - **View Details**— See Viewing the iSCSI ports details on the next page.

Viewing the iSCSI ports details

This procedure explains how to view the iSCSI ports details.

1. Select a storage system running HYPERMAX 5977 or higher.
2. Select **System › ISCSI Dashboard** to open the iSCSI dashboard.
3. Click **iSCSI Ports** to open the **iSCSI Ports** list view and select a specific director to filter the list.
4. Select a row from the list and click **View Details** to open the **iSCSI Ports** details view.

Properties panel

The following properties display, depending on the operating environment:

- **Dir: Port**— The name of the director and the iSCSI port combined.
- **Status** — The current status of the iSCSI port (offline or online).
- **Number of IP Interfaces**— The total number of IP interfaces attached to this iSCSI port.
- **Number of iSCSI Targets**— The total number of iSCSI targets to all the IP interfaces on the selected iSCSI port.
- **Volumes**—The total number of volumes mapped to all iSCSI targets on this port.
- **Speed GB/Sec**— The speed of the iSCSI port.

Related Objects panel

The following objects display, depending on the operating environment:

- **IP Interfaces** — The list view of IP interfaces that are associated with the physical ports on this director.
- **Attached iSCSI Targets** —The list view of iSCSI targets on this director that are attached to IP interfaces.
- **Unattached iSCSI Targets** — The list view of iSCSI targets on this director that are unattached.
- **Mapped Volumes** — The list view of volumes that are mapped to iSCSI targets associated with the director.
- **iSCSI Director** — The list view listing the director that is associated with the selected iSCSI port .